

## IMPACT OF DIFFERENT TRAJECTORIES ON EXTRINSIC SELF-CALIBRATION FOR VEHICLE-BASED MOBILE LASER SCANNING SYSTEMS

Markus Hillemann<sup>1,2\*</sup>, Jochen Meidow<sup>2</sup>, Boris Jutzi<sup>1</sup>

<sup>1</sup> Institute of Photogrammetry and Remote Sensing, Karlsruhe Institute of Technology, Karlsruhe, Germany - markus.hillemann@kit.edu

<sup>2</sup> Fraunhofer Institute of Optronics, System Technologies and Image Exploitation IOSB, Ettlingen, Germany

**KEY WORDS:** Self-Calibration, Mobile Laser Scanning, Relative Orientation, Boresight & Lever Arm

### ABSTRACT:

The extrinsic calibration of a Mobile Laser Scanning system aims to determine the relative orientation between a laser scanner and a sensor that estimates the exterior orientation of the sensor system. The relative orientation is one component that limits the accuracy of a 3D point cloud which is captured with a Mobile Laser Scanning system. The most efficient way to determine the relative orientation of a Mobile Laser Scanning system is using a self-calibration approach as this avoids the need to perform an additional calibration beforehand. Instead, the system can be calibrated automatically during data acquisition. The entropy-based self-calibration fits into this category and is utilized in this contribution. In this contribution, we analyze the impact of four different trajectories on the result of the entropy-based self-calibration, namely (i) uni-directional, (ii) ortho-directional, (iii) bi-directional, and (iv) multi-directional trajectory. Theoretical considerations are supported by experiments performed with the publicly available *MLS 1 – TUM City Campus* data set. The investigations show that strong variations of the yaw angle in a confined space or bidirectional trajectories as well as the variation of the height of the laser scanner are beneficial for calibration.

### 1. INTRODUCTION

To obtain a detailed 3D representation of the environment, Laser Scanning, also known as Light Detection And Ranging (LiDAR), is a very effective and accurate technique. Depending on the task, it can be applied on different scale levels. Airborne Laser Scanning, for example, often is used to create a Digital Terrain Model and a Digital Surface Model, respectively, on a regional scale. On a local scale, besides classical Terrestrial Laser Scanning, Mobile Laser Scanning (MLS) is increasingly becoming the focus of practical application, since it is efficient and provides comparable accuracy. Laser Scanning is, for instance, used to generate 3D city models, building models or models of indoor environments. For this purpose, different mobile platforms like vehicles, Unmanned Aerial Vehicles (UAVs), or even smaller robots to navigate indoors can be utilized.

The relative orientation between a laser scanner and a sensor that estimates the exterior orientation, also known as pose, of the sensor system is one component that limits the accuracy of a 3D point cloud which is captured with an MLS system (Skaloud & Lichti, 2006; Habib et al., 2009; Jutzi et al., 2014). The process of estimating the relative orientation between two sensors is known as extrinsic calibration. The six parameters of the relative orientation can be divided into three parameters of the relative translation (a.k.a. lever arm) and three parameters of the relative rotation (a.k.a. boresight) between the two sensors. They can also be interpreted as a spatial transformation, i.e. a 3D motion, between the coordinate frame of the laser scanner and the coordinate frame of the navigation unit.

In contrast to frequently used semi-automatic calibration approaches which rely on calibration targets (Geiger et al., 2012; Vasconcelos et al., 2012; Heinz et al., 2017; Urban

& Jutzi, 2017; Hillemann & Jutzi, 2017), a self-calibration approach (Habib et al., 2011; Sheehan et al., 2012) is the most practical way of determining the relative orientation of a Mobile Laser Scanning system as this avoids the need for performing an additional calibration procedure beforehand. The entropy-based self-calibration (Sheehan et al., 2012) fits into this category and is used for the experiments. The approach has been used to calibrate different vehicle-based MLS systems (Maddern et al., 2012, 2017).

The entropy-based self-calibration does not consider assumptions about the environment except that it is static. However, there are general constraints about the vehicle trajectory. It is recommended to include variation in all six degrees of freedom to effectively apply the parameter estimation (Maddern et al., 2012). As consequence, the trajectory includes translation in three perpendicular directions (across driving direction, along driving direction and vertical) and rotation about three perpendicular axes (roll, pitch and yaw). For vehicle-based MLS systems large horizontal translations and a varying yaw angle can be realized. However, the vertical translations and the variation of the roll and the pitch angle are limited and generally very small. This raises the question of whether an accurate estimation of all calibration parameters is even possible for vehicle-based MLS.

This contribution is organized as follows. To analyze the impact of the trajectory on the calibration of a MLS system, we present theoretical considerations in Section 2. In Section 3 the entropy-based self-calibration is stated. To support the theoretical considerations we perform experiments on real data in Section 4. The real data we use are excerpts from the publicly available data set *MLS 1 – TUM City Campus* (Gehring et al., 2017). These excerpts represent four different trajectories driven in the city of Munich. In Section 5 we discuss the results and Section 6 concludes this contribution.

\*Corresponding author

## 2. IMPACT OF THE TRAJECTORY

As mentioned, to determine the parameters of the relative orientation, the trajectory must include translation and rotation in all six degrees of freedom (Maddern et al., 2012). The research of Habib et al. (2009) and Ravi et al. (2018) deal in more detail with the question concerning the impact of the trajectory on the result of the calibration in the context of airborne laser scanning and UAV-based laser scanning, respectively.

Further, Ravi et al. (2018) try to find a minimal flight configuration for accurately calibrating a UAV-based MLS system based on overlapping strips. The authors' considerations and findings are also interesting for vehicle-based MLS or comparable calibration approaches. They study the impact of each parameter of the relative orientation on the point cloud. For better interpretation the coordinate frames of the laser scanner and the navigation system are aligned. In the following the key findings regarding the relative translation are reflected:

- The relative translation *across motion direction* causes the points to shift across motion direction. To make the impact of this parameter visible, Ravi et al. (2018) suggest to scan vertical objects parallel to motion direction from opposite motion directions.
- The relative translation *along motion direction* causes the points to shift along driving direction. To constrain this parameter the authors suggest to scan vertical objects perpendicular to motion direction from opposite motion directions.
- The relative translation in *vertical direction* causes the points to shift in vertical direction. In practice, the vertical relative translation component cannot be calibrated.

Besides the relative translations the relative rotations are of interest. Each parameter of the relative rotation causes the points to shift in two spatial directions:

- Changing the *roll angle* of the relative rotation causes the points to shift across motion direction and additionally in vertical direction. To estimate relative roll the authors suggest to scan objects with varying height w.r.t. the height the laser scanner moves in.
- Analogously, changing the *pitch angle* causes the points to shift along motion direction and additionally in vertical direction. Again, to estimate relative pitch the authors suggest to scan objects with varying height w.r.t. the height the laser scanner moves in.
- Changing the *yaw angle* causes the points to shift across and along motion direction. The suggested solution is to scan at least two vertical objects with considerable lateral separation from two parallel trajectories with the same motion direction.

In summary, a minimal motion configuration to effectively constrain the calibration parameters is the following: It is recommended that the MLS scans the same three-dimensional objects with different vertical distances between the objects and the laser scanner. Additionally, it is recommended to use at least three parallel trajectories, two with the same motion direction and one with the opposite motion direction.

## 3. METHODOLOGICAL BACKGROUND OF THE ENTROPY-BASED SELF-CALIBRATION

In this section the entropy-based self-calibration is introduced. Therefore, we explain the relation between the calibration and registration especially in the context of point cloud processing. Further, we complete the calibration methodology by specifying the optimization strategy.

### 3.1 Relation between Calibration and Registration

The calibration of a laser scanning system is closely related to the task of point cloud registration. Finding spatial transformations which align various point clouds implies that the MLS system is calibrated, i.e., the poses and the relative orientation parameters of the MLS system are known. Conversely, if the point clouds are aligned, we conclude that the relative orientation of the MLS system is found. In the following, we consider only rigid transformations for the mapping between point clouds, consisting of translations and rotations.

Methods to solve the point cloud registration task include especially the *Iterative closest point* algorithm (Besl & McKay, 1992), *Robust point matching* (Gold et al., 1998), *Coherent point drift* (Myronenko & Song, 2010), and *Kernel Correlation* (Tsin & Kanade, 2004). Compared with the Iterative closest point (ICP) algorithm, the Kernel Correlation approach is more robust against noisy data, since within this methods – unlike the ICP algorithm – every scene point affects every model point. Thus, the approach exploits multiple links for each point, yields more often correct registration results, and therefore serves as a reference in the following.

Kernel correlation extends the concept of correlation to point sets. The correlation value itself is considered as a measure of affinity or similarity of two point clouds or as a measure of compactness for the merged cloud (Tsin & Kanade, 2004). Given a point cloud  $\mathcal{X} = \{\chi_i\}$ ,  $i = 1, \dots, N$  with  $N$  points  $\chi_i$ , the kernel correlation

$$KC(\mathcal{X}) = \sum_i^K KC(\chi_i, \mathcal{X} \setminus \chi_i) \quad (1)$$

is the sum of the kernel correlations for all points w.r.t. all other points. If the points in the point cloud are close to each other, the KC-value (1) is large. The kernel correlation between a point  $\chi_i$  and the point cloud  $\mathcal{X}$  is in turn

$$KC(\chi_i, \mathcal{X} \setminus \chi_i) = \sum_{i \neq j} K(\mathbf{x}_i, \mathbf{x}_j) \quad (2)$$

with a kernel function  $K$ , e.g., a Gaussian kernel

$$G(\mathbf{x}_i, \mathbf{x}_j, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{D/2}} \exp\left\{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}\right\} \quad (3)$$

with the free parameter  $\sigma$  and  $D$  denoting the dimension of the vector with coordinates, in our case  $D = 3$ . Thus, the kernel is a function of the Euclidean distance between the two points  $\chi_i$  and  $\chi_j$ .

Using the kernel (3) and a Gaussian mixture model, a non-parametric density  $p(\mathcal{X})$  of a point cloud  $\mathcal{X}$  can be specified by a kernel density estimation  $p(\chi_i)$  at each point

position. This allows establishing a simple relationship to the so-called collision, correlation or Rényi's quadratic entropy (RQE), as a special case of the Rényi entropy (Rényi, 1961)

$$H_2(\mathcal{X}) = -\log \sum_{i=1}^N (p(\mathcal{x}_i))^2 \quad (4)$$

which can be used to specify the compactness of a mixture density. The measure computes the negative logarithm of the probability that two independent samples from the distribution are the same. Thus, the relation

$$KC(\mathcal{X}) \propto C + \exp(-H_2(\mathcal{X})) \quad (5)$$

holds with the constant  $C$  (Tsin & Kanade, 2004).

finding the minimum entropy configuration of the joint point cloud.

### 3.2 Estimation of the Relative Orientation by Minimizing the Entropy

The entropy-based self-calibration approach can be interpreted as minimizing a point cloud quality error to optimize approximate parameters of the relative orientation between the laser scanner and the navigation unit. The calibration is formulated as a non-linear least squares optimization problem. The cost function  $f$  of this problem is based on the Rényi quadratic entropy (RQE) and is defined as

$$f = -\sum_{i=1}^N \sum_{j=1}^N G(\mathbf{x}_i, \mathbf{x}_j, \sigma^2), \quad (6)$$

where  $N$  is the number of points in the point cloud,  $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^3$  are 3D points, and  $G(\mathbf{x}_i, \mathbf{x}_j, \sigma^2)$  is the Gaussian function with variance  $\sigma^2$ .

By using an isotropic kernel and the same variance  $\sigma^2$  for each dimension, the parameters of the cost function are reduced to a single parameter  $\sigma$  (Maddern et al., 2012). The smaller  $\sigma$  is, the peakier is the Gaussian function.  $\sigma$  can also be interpreted as prior knowledge of the precision of a captured 3D point. The entropy-based self-calibration uses a fixed value  $\sigma$  for the cost function and consequently for the entire point cloud.

## 4. EXPERIMENTS

As mentioned, we perform experiments on the publicly available *MLS 1 – TUM City Campus* data set (Gehring et al., 2017). We extrinsically calibrate one of the 3D laser scanners *Velodyne HDL-64E* with respect to the navigation system *Applanix POS LV*. Both sensors are mounted on top of a moving vehicle. The *Applanix POS LV* fuses the measurements of GNSS receivers and an Inertial Measurement Unit (IMU) hence delivering 6-DOF poses with a rate of up to 200 Hz. The considered laser scanner *Velodyne HDL-64E* is positioned on the front roof of the vehicle with an angle of  $25^\circ$  to the horizontal plane and a slope to the front with an angle of  $45^\circ$ . The laser scanner has a measurement range of 120 m and a range measurement accuracy of 5 cm ( $1\sigma$ ). It has a vertical field of view of  $26.8^\circ$  and a horizontal field of view of  $360^\circ$ . The ground truth for the relative orientation between the 3D laser scanner and the navigation system is provided with the data set.

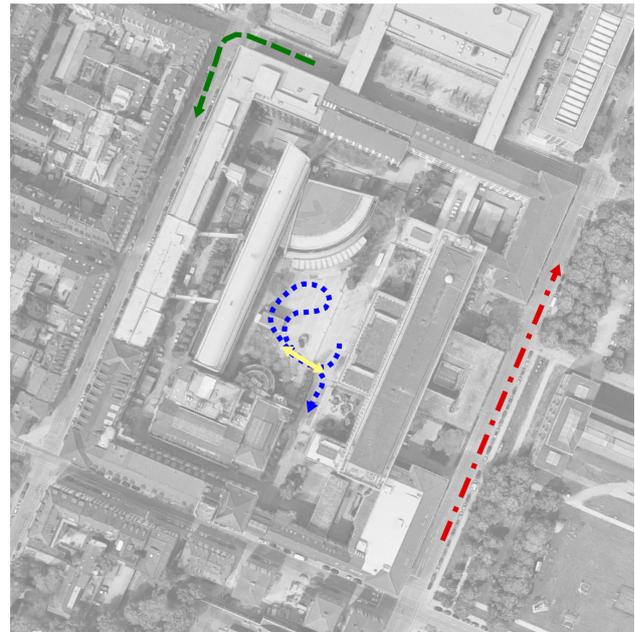


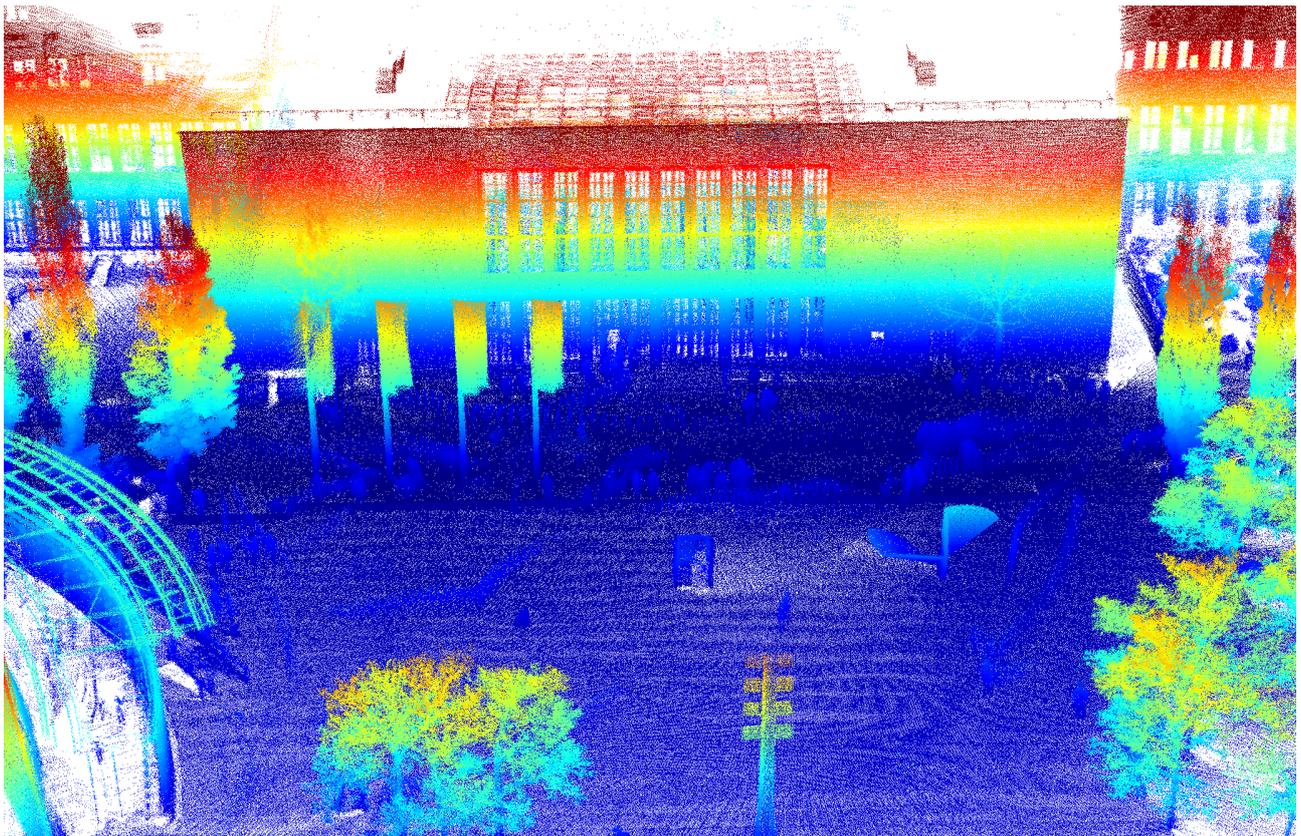
Figure 1. Overview of the trajectories. Uni-directional trajectory (dash-dotted red line), ortho-directional trajectory (dashed green line), bi-directional trajectory (solid yellow line), multi-directional trajectory (dotted blue line).

We determine the relative orientation four times each time based on a different trajectory. The trajectories are depicted in Figure 1. The first one represents a *uni-directional* trajectory (dash-dotted red line) like driving straight along the road. We call the second one *ortho-directional* trajectory (dashed green line). This trajectory is typical for turning a right-angled corner. The third one represents a *bi-directional* trajectory (solid yellow line) like driving the same road back and forth. The fourth one is a *multi-directional* trajectory (dotted blue line) like for turning around at a confined space. The bi-directional and the multi-directional trajectory, moreover, contain a ramp so that the height of the laser scanner w.r.t. the ground changes. All trajectories with exception of the bi-directional trajectory have the same number of poses. The bi-directional trajectory consists of a third of the poses, since the data set does not contain a longer sequence for a trajectory like this.

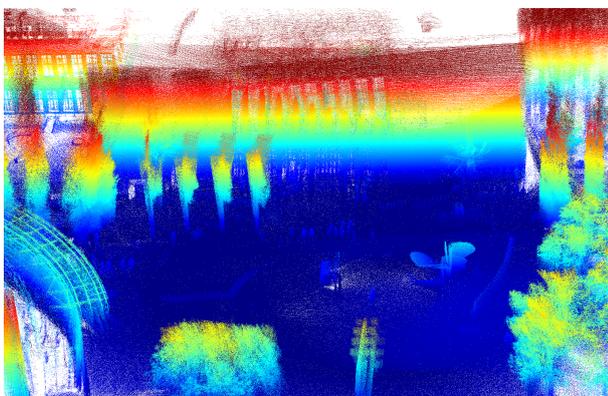
The length of the trajectories varies due to the different driving speeds. The approximate parameters of the relative orientation for optimization have an error of +1 m for each relative translation and  $+3^\circ$  for each relative rotation angle. To evaluate the precision of the entropy-based self-calibration for the different trajectories, we randomly sub-sample the point cloud twenty times and compute one calibration result per sample.

As the impact of changing the relative rotation parameters on the point cloud is much higher than changing the relative translation parameters, we perform the calibration in three steps:

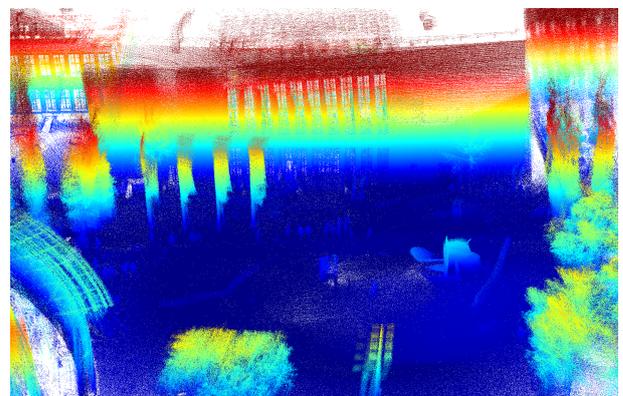
1. Coarse estimation of the relative rotation parameters.
2. Coarse estimation of the relative translation parameters with the improved relative rotation parameters.
3. Optimization of all parameters of the relative orientation.



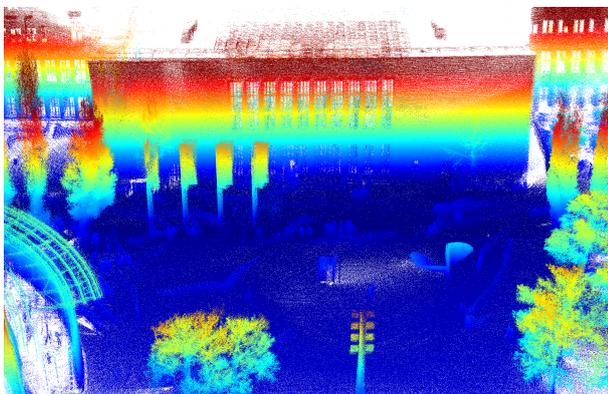
(a)



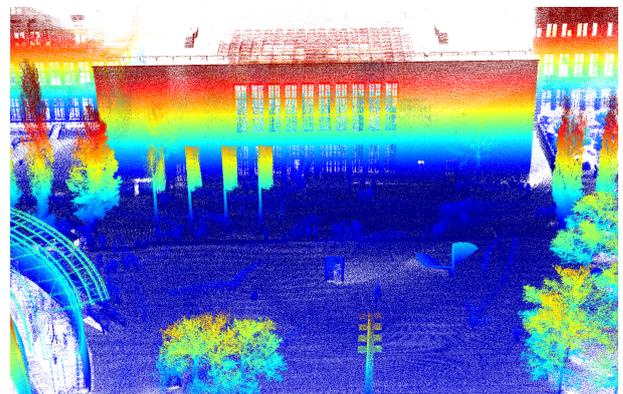
(b)



(c)



(d)



(e)

Figure 2. Point clouds computed after calibration with different trajectories. (a) Ground truth calibration, (b) uni-directional trajectory, (c) ortho-directional trajectory, (d) bi-directional trajectory, (e) multi-directional trajectory.

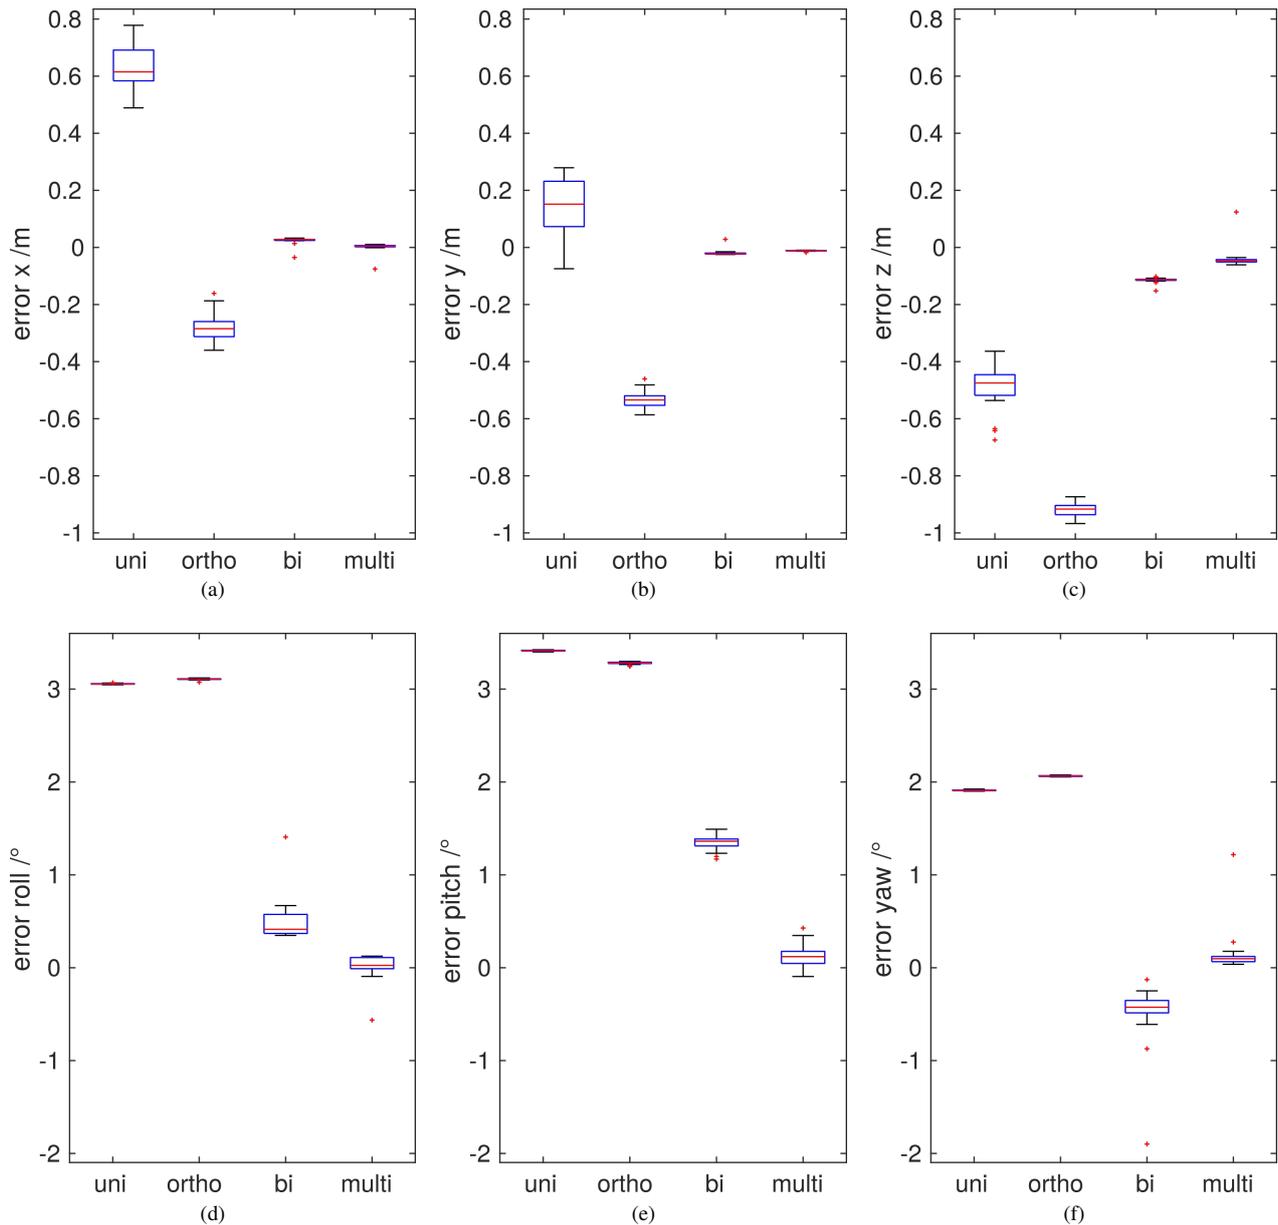


Figure 3. Calibration results for the relative translation parameters (a-c) and the relative rotation parameters (d-f) based on the four different investigated trajectories. The approximate parameters of the relative orientation before calibration are  $x_0 = y_0 = z_0 = 1\text{ m}$  and  $\text{roll}_0 = \text{pitch}_0 = \text{yaw}_0 = 3^\circ$ .

Figure 2 illustrates the results of the calibrations qualitatively. The figure shows point clouds of a courtyard computed with the different relative orientations. Figure 2a represents the ground truth calibration, Figure 2b the uni-directional, Figure 2c the ortho-directional, Figure 2d the bi-directional, and Figure 2e the multi-directional trajectory. The point clouds in Figure 2b and 2c appear distorted. For example the windows of the buildings are hardly visible. However, Figure 2d and 2e show that the entropy-based self-calibration achieves very similar results compared to the ground truth. Qualitative differences between these three results are visually indiscernible.

Figure 3 shows the calibration results for the four different trajectories. For all trajectories the entropy-based self-calibration improves the approximate relative translations. However, the medians of the errors for the uni-directional trajectory and the ortho-directional trajectory are in a range

of decimeters. The errors for the bi- and the multi-directional trajectory, on the other hand, are in a range of centimeters. The results of the relative rotations for the uni- and the ortho-directional trajectories show an error of  $3^\circ$  after calibration. At the same time, the result for these trajectories varies by less than  $0.1^\circ$ . The result of the bi-directional trajectory shows smaller errors than the first two trajectories, however, the error of the pitch angle is larger than  $1^\circ$ . The multi-directional trajectory leads to the smallest errors.

## 5. DISCUSSION

The qualitative as well as the quantitative results show that the trajectory has a major impact on the calibration result. For all trajectories the approximate parameters of the relative translation and for the relative yaw angle are improved, however, for the uni-directional and the ortho-directional

trajectory the result is not accurate. The error of the relative roll and pitch angle even degrades after calibration. So these two trajectories are unsuitable for calibration.

However, the precision of the estimation of the relative roll and pitch angle is very high for the uni-directional and the ortho-directional trajectory. Consequently, these results show that it is not sufficient to use only the precision as an evaluation metric. In our case, there is an independent measurement of the parameters of relative orientation for comparing the results. In the absence of ground truth, it is recommended to consider additional evaluation metrics such as accuracy at control points (Csanyi & Toth, 2007) or average point to plane distances in planar regions of the point cloud (Skaloud & Lichti, 2006).

The bi-directional trajectory is close to the minimum requirements for trajectories given in Section 2. The bi-directional trajectory leads to a considerably better result than the uni- and the ortho-directional trajectory. Since the number of poses and thus the number of points for this trajectory is the lowest, the calculation with this trajectory is the fastest. The multi-directional trajectory leads to the best results. Especially the error of the pitch angle is even smaller for the multi-directional trajectory. The difference between the bi-directional and the multi-directional trajectory is primarily that the multi-directional trajectory contains even more variations of the yaw angle. This is supposed to be the reason why the multi-directional trajectory yields the best result. The result with the multi-directional trajectory and the ground truth is visually indiscernible. In addition, each parameter can be estimated with satisfactory accuracy, although the underlying motion changes considerably in only four of the six possible degrees of freedom.

It is important to note that the vertical relative translation can also be estimated for the bi-directional and the multi-directional trajectory, although with the worst precision of the three components of relative translation. One reason (i) is the configuration of the considered MLS system. As mentioned, the laser scanner is mounted at an angle of  $25^\circ$  to the horizontal and a slope to the front with an angle of  $45^\circ$ . The inaccurate estimation in vertical direction is distributed among the other parameters. Further, the different optimization strategies of the entropy-based self-calibration supposed to be another reason (ii) that the vertical relative translation can also be estimated with this approach. The calibration procedure used by Ravi et al. (2018) minimizes the variation of points at planar structures along the surface normal. Therefore, a point at a planar structure scanned multiple times from different viewing angles (different trajectories, resp.) only provides information in one direction for the determination of the parameters. In contrast, when taking edges into account and minimizing the entropy of the point cloud, each point at an edge provides information in two perpendicular directions. Further taking corners into account even provides information in three perpendicular directions. However, typically for laser scanning the range measurements are more inaccurate for points at edges or corners. Consequently, there is a trade off for points at structures like these.

## 6. CONCLUSION AND OUTLOOK

We analyze the impact of the trajectory on the extrinsic calibration result based on reflecting theoretical considerations and supporting these with real data experiments. The

experiments show that the underlying trajectory has a large impact on the calibration result. It is insufficient for the entropy-based self-calibration to scan the same part of the environment with only one motion direction even if the trajectory contains curves. The trajectory requirements that apply to the calibration of UAV-based MLS systems (Ravi et al., 2018) can be confirmed for the calibration of vehicle-based MLS systems. A trajectory with two opposite directions is useful for determining the parameters of the relative orientation. The same applies to the change of the height of the sensor w.r.t. the environment. Additionally, strong variations of the yaw angle in a confined space, i.e., tight curves, are beneficial to increase the variation of the angle of view on the same parts of the environment. Further, the investigations show that it is not sufficient to use only the precision of the procedure as an evaluation metric.

Future work consists of the development of an automatic pose selection algorithm to effectively constrain the parameter estimation and to enable accurate online calibration.

## References

- Besl, P., McKay, N., 1992. A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14, 239–256.
- Csanyi, N., Toth, C. K., 2007. Improvement of lidar data accuracy using lidar-specific ground targets. *Photogrammetric Engineering & Remote Sensing*, 73, 385–396.
- Gehring, J., Hebel, M., Arens, M., Stilla, U., 2017. An Approach to Extract Moving Objects from MLS Data Using a Volumetric Background Representation. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-1, 107–114.
- Geiger, A., Moosmann, F., Car, O., Schuster, B., 2012. Automatic camera and range sensor calibration using a single shot. *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 3936–3943.
- Gold, S., Rangarajan, A., Lu, C., Suguna, P., Mjolsness, E., 1998. New algorithms for 2d and 3d point matching: pose estimation and correspondence. *Pattern Recognition*, 38, 1019–1031.
- Habib, A. F., Kersting, A. P., Bang, K. I., Lee, D., 2009. Alternative methodologies for the internal quality control of parallel LiDAR strips. *IEEE Transactions on Geoscience and Remote Sensing*, 48, 221–236.
- Habib, A. F., Kersting, A. P., Shaker, A., Yan, W., 2011. Geometric calibration and radiometric correction of LiDAR data and their impact on the quality of derived products. *Sensors*, 11-9, 9069–9097.
- Heinz, E., Eling, C., Wieland, M., Klingbeil, L., Kuhlmann, H., 2017. Analysis of different reference plane setups for the calibration of a mobile laser scanning system. In: Lienhart, W. (Hrsg.): *Ingenieurvermessung 17, Beiträge zum 18. Internationalen Ingenieurvermessungskurs, Graz, Austria*, 131–146.
- Hillemann, M., Jutzi, B., 2017. UCalMiCeL—Unified intrinsic and extrinsic calibration of a Multi-Camera-System and a Laserscanner. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, IV-2, 17–24.

Jutzi, B., Weinmann, M., Meidow, J., 2014. Weighted Data Fusion for UAV-borne 3D Mapping with Camera and Line Laser Scanner. *International Journal of Image and Data Fusion*, 5, 226–243.

Maddern, W., Harrison, A., Newman, P., 2012. Lost in translation (and rotation): Rapid extrinsic calibration for 2D and 3D LIDARs. *Proceedings of the 2012 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 3096–3102.

Maddern, W., Pascoe, G., Linegar, C., Newman, P., 2017. 1 year, 1000 km: The Oxford RobotCar dataset. *The International Journal of Robotics Research*, 36, 3–15.

Myronenko, A., Song, X., 2010. Point set registration: Coherent Point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32, 2262–2275.

Ravi, R., Shamseldin, T., Elbahnasawy, M., Lin, Y., Habib, A., 2018. Bias Impact Analysis and Calibration of UAV-Based Mobile LiDAR System with Spinning Multi-Beam Laser Scanner. *Applied Sciences*, 8, 297.

Rényi, A., 1961. On measures of entropy and information. *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability*, 1, University of California Press, 547–561.

Sheehan, M., Harrison, A., Newman, P., 2012. Self-Calibration for a 3D Laser. *The International Journal of Robotics Research*, 31, 675–687.

Skaloud, J., Lichti, D., 2006. Rigorous approach to bore-sight self-calibration in airborne laser scanning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 61–1, 47–59.

Tsin, Y., Kanade, T., 2004. A correlation-based approach to robust point set registration. *Computer Vision ECCV*, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 558–569.

Urban, S., Jutzi, B., 2017. LaFiDa – A Laserscanner Multi-Fisheye Camera Dataset. *Journal of Imaging*, 3, 5. <http://www.mdpi.com/2313-433X/3/1/5>.

Vasconcelos, F., Barreto, J.P., Nunes, U., 2012. A Minimal Solution for the Extrinsic Calibration of a Camera and a Laser-Range-finder. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 34, 2097–2107.