

Maximum Likelihood Remission Calibration for Groups of Heterogeneous Laser Scanners

Bastian Steder

Michael Ruhnke

Rainer Kümmerle

Wolfram Burgard

Abstract—Laser range scanners are commonly used in mobile robotics to enable a robot to sense the spatial configuration of its environment. In addition to the range measurements, most scanners provide remission values, representing the intensity of the returned light pulse. These values add a visual component to the measurement and can be used to improve reasoning on the data. Unfortunately, a remission value does not directly tell us how bright a measured surface is in the infrared spectrum. Rather, it varies with respect to the incidence angle and the range at which it was measured. In addition, multiple scanners typically do not agree upon the values of a certain surface. In this paper, we present a calibration method for remission values of multiple laser scanners considering dependencies in range, incidence angle of the measured surface, and the respective scanner unit. Our system learns the calibration parameters based on a set of registered point clouds. It uses a graph optimization scheme to minimize the error between different measurements, so that all involved scanners yield consistent reflection values, independent of the perspective from which the corresponding surface is observed.

I. INTRODUCTION

Laser range scanners are widely used in mobile robotics since they provide highly accurate spatial information of the vicinity of the robot. Typical applications in this context are mapping, localization, obstacle avoidance, modeling, or recognition tasks. The underlying principle of laser range scanners is to send out focused infrared light and measure the time until the reflected light returns to the sensor to determine the distance to the surface. Besides the measured range, most laser range scanners also return remission values, also called intensity or reflectivity values. These values correspond to the amount of emitted light that returned to the measuring unit and can tell us how bright an object appears in the infrared spectrum used by the sensor. These measurements add a visual appearance component to the spatial information of the laser scanner and can be used to improve localization [7], extrinsic calibration to visual sensors [1], or classification tasks [13]. Compared to traditional visual camera images, remission values have the advantage of being independent to external lighting conditions in typical applications. Furthermore, no extrinsic calibration to the reference frame of the laser scanner is needed since we measure them at the same time with the same device.

Unfortunately, the measured remission values are not purely dependent on the reflectivity of the measured material,

B. Steder, M. Ruhnke, and W. Burgard are members of the University of Freiburg, Department of Computer Science, D-79110 Freiburg, Germany. R. Kümmerle is with KUKA Roboter GmbH, Augsburg, Germany.

This work has been partly supported by the EC under the grant numbers ERC-AG-PE7-267686-LifeNav, FP7-610603-EUROPA2, and H2020-645403-RobDREAM.

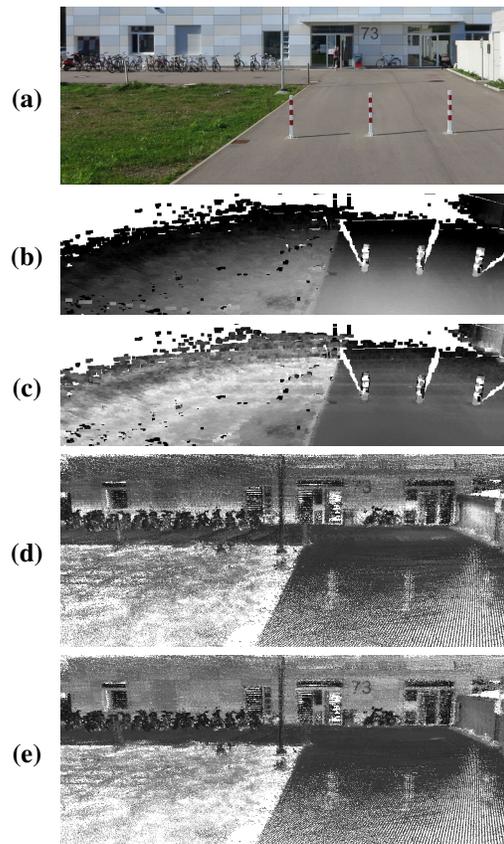


Fig. 1. Example Scene, corresponding scans and maps obtained with calibrated and uncalibrated remission values: (a) Observed scene. (b) Uncalibrated 3D point cloud of one scan from a single tilting Hokuyo UTM laser scanner. The measurements at the bottom center are brighter, because of closer ranges and steeper incidence angles. (c) Calibrated version of the above. The road shows continuous gray values now. (d) Uncalibrated 3D point cloud captured using a combination of three different 3D scanners, a Velodyne HDL-32E, a downwards facing SICK LMS 151, and the tilting Hokuyo UTM from (b)&(c). The structures are not clearly visible and remission values on continuous surfaces are not uniform. (e) Calibrated version of (d), showing better contrast and more uniform gray values.

but also on other factors like the range at which it was observed, the incidence angle of the measured surface and the brightness of the specific laser scanner used. See Fig. 1 for example images.

Multi-beam lasers, like the Velodyne range scanners, should yield consistent intensity values of a certain point on the measured surface to provide an accurate view of the true reflectivity attributes of the observed environment. Furthermore, it seems desirable to calibrate the intensity values reported by all laser range scanners on a mobile robot. This is an important prerequisite to perform a classification

on a fused point cloud with the same classifier or to build a consistent high resolution 3D model with laser intensity texture.

In this paper, we present a novel method to learn how the range, incidence angle and the properties of an individual laser scanner influence the measurements and thereby provide a calibration procedure to enable a system to better estimate the reflective properties of an observed structure without these influences. Our method does not rely on any special calibration structures or setting. Instead it uses a collection of registered 3D point clouds, captured by one or more laser scanners on a moving robot. If the same spot in the environment is observed by possibly different scanners from different perspectives and distances, it is the goal to learn calibration parameters that, when applied to the measurements, lead to consistent measurements of the reflectivity in that spot. Our system uses a least square error minimization technique based on a graph built from corresponding observations to learn the calibration parameters in an unsupervised fashion. Our method directly accounts for the dependence of measured range, incidence angle and brightness of the specific laser scanner.

The system described in this paper will be made available as open source software [11].

II. RELATED WORK

In the past, several researchers addressed the problem of calibrating multi laser setups like the Velodyne HDL or custom made setups [6], [8], [9]. Those approaches, however, focus on the extrinsic parameters and also on the time skew of the lasers but do not account for calibrating the intensity data of different lasers. Such intensity data has been proven as useful, e.g., for localizing a vehicle in urban environments [4] or classifying the terrain [13].

The closest to our work is the approach by Levinson and Thrun [5]. They describe an unsupervised approach for extrinsic and intrinsic calibration of multi-beam laser scanners. As part of their work the authors present an expectation maximization calibration method for intensity values. The result of this method is a lookup table for every different laser in the beam array that maps a measured intensity to the most likely intensity value. In contrast, our method models the relation between the range as well as the incidence angle and the perceived intensity value in a global optimization system to find the maximum likelihood calibration. Note that also other researchers reported a crucial impact of the incidence angle on both the range and the intensity [10], [13].

In the context of air vehicles equipped with a laser Kaasalainen *et al.* [2] proposed a calibration scheme for laser intensity values. The authors explicitly model the dependencies between intensity measurement and the range and incidence angle at which the intensity was measured. Their approach, however, requires specific calibration patterns with known reflectance attributes, whereas our method can be applied to an arbitrary dataset with sufficient variation in viewpoints from which the individual surfaces in the environment were observed.

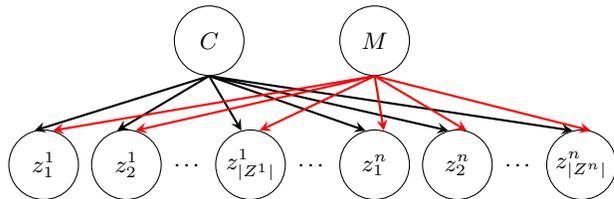


Fig. 2. This figure shows the Bayesian network of the problem we want to address in this paper. Each measurement z_k^l , which is the k -th measurement of the laser with the index l , depends on the reflectivity of a certain cell $m \in M$ and on the calibration parameters C . It is our goal to determine the hidden variables C and M .

III. SYSTEM DESCRIPTION

In this section, we describe the underlying calibration problem, discuss different independence assumptions as well as how we include those in our model, and elaborate how we formulate this as a graph optimization problem which we iteratively solve using least squares.

A. Problem Definition

The input to our system is a set of registered 3D point clouds (meaning we can perform mapping with known poses), captured with a set of one or more laser range scanners. The point clouds have to include the measured remission value for every 3D point and also the sensor position from which the point was observed. Given such data, it is our goal to determine a calibration function C that normalizes the measured remission values such that the values are consistent over the whole dataset. More formally: Let $z_k^l \in Z^l$ be the k -th measurement captured by the laser with index $l \in L$. Each of these measurements contains a triple of values $z_k^l = (i_k^l, r_k^l, \alpha_k^l)$, where i_k^l is the measured remission/intensity value, r_k^l is the measured range value, and α_k^l is the incidence angle at which the laser beam hits the surface. We added the latter to the measurement for simplicity. The laser does not directly measure the incidence angle but we can estimate it from a local neighborhood around the measured point. Please note that multi-beam laser scanners, like those produced by Velodyne, contribute to the set of laser indices L with every individual laser beam they contain.

It is our goal to calculate the calibrated remission value i_k^l for such a measurement:

$$i_k^l = C(l, z_k^l). \quad (1)$$

Fig. 2 visualizes the Bayes Network for the measurements. The calibration function C is supposed to normalize the measured remission values in such a way that measurements from different perspectives and from different lasers yield the same value for a certain position in the map M . Assuming Gaussian noise on the remission measurements, the optimal guess for the reflectivity of a map cell $m \in M$ is the mean of the calibrated remission values.

In the next section we will present concrete implementations of the calibration function C .

B. Calibration Parameters

In our system, we implemented different concrete definitions of C that represent different independence assumptions. The most general calibration function that we tested, with no underlying independence assumptions is

$$C^1(l, z_k^l) = c^1(r_k^l, \alpha_k^l, l) \cdot i_k^l, \quad (2)$$

where $c^1 \in \mathbb{R}^+$ is a factor for the measured remission value.

Under the assumption that the influences of the range and of the incidence angle are mutually independent given the laser index, we can choose an alternative calibration function:

$$C^2(l, z_k^l) = c_r^2(r_k^l, l) \cdot c_\alpha^2(\alpha_k^l, l) \cdot i_k^l, \quad (3)$$

where $c_r^2 \in \mathbb{R}^+$ and $c_\alpha^2 \in \mathbb{R}^+$ are individual factors for the range and incidence angle.

The functions C^1 and C^2 both contain different calibration parameters for every individual laser. We can go one step further and add the assumption that lasers of the same type (e.g., multiple SICK LMS 151, multiple Hokuyo UTM, or the different beams inside of a Velodyne scanner) behave similarly regarding the measurement of remission values. Based on this, we obtain

$$C^3(l, z_k^l) = c_r^3(r_k^l, T(l)) \cdot c_\alpha^3(\alpha_k^l, T(l)) \cdot c_i^3(l) \cdot i_k^l, \quad (4)$$

where $c_r^3 \in \mathbb{R}^+$, $c_\alpha^3 \in \mathbb{R}^+$, and $c_i^3 \in \mathbb{R}^+$ are individual factors for the range, incidence angle, and the laser index. The term $T(l) \in \mathbb{N}^+$ is a type identifier, which assigns the same number to laser scanners of the same type.

C. Formulation as an Optimization Problem

We can obtain the maximum a posteriori estimate for the calibration parameters and the map of reflectivity values as

$$(C, M)^* = \operatorname{argmax}_{C, M} p(C, M | Z, L) \quad (5)$$

$$= \operatorname{argmin}_{C, M} -\log p(C, M | Z, L), \quad (6)$$

where $p(\cdot)$ is a probability distribution, C is the set of calibration parameters, $M = \{m_1, \dots, m_{|M|}\}$ is the map containing individual cells each with an associated reflectivity value, Z is the set of all sensor measurements and L is the set of individual laser scanner indices.

If we assume that the parameters are independent and $p(\cdot)$ is a Gaussian, we can rewrite Eq. (6) as

$$(C, M)^* = \operatorname{argmin}_{C, M} \sum_{l \in L, z_k^l \in Z} \|e(z_k^l, M)\|_{\Sigma_k^l}^2, \quad (7)$$

where $\|e\|_{\Sigma}^2 = e^T \Sigma^{-1} e$ is the squared Mahalanobis distance, which here represents the exponential part of the Gaussian $p(\cdot)$. We solve (7) with the g^2o framework [3].

For each cell $m_j \in M$ the individual lasers should measure the same reflectivity value. Consequently, the error function $e(\cdot)$ computes the difference between the calibrated reflectivity value $C(l, z_k^l)$ of the laser beam z_k^l falling into the cell m_j and the estimated reflectivity value:

$$e(z_k^l, M) = C(l, z_k^l) - m_j. \quad (8)$$

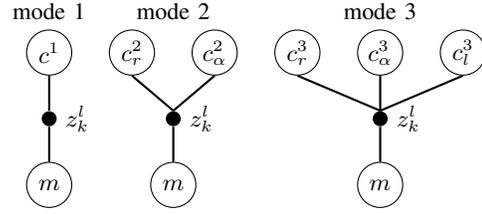


Fig. 3. Graph structure for a single constraint for the three different calibration modes.

D. Optimization Graph

The vertices in the graph are the individual calibration factors c , as described in Sec. III-B. Each calibration factor c relates to a laser generating a range measurement with a specific incidence angle. As the number of nodes in the graph has to be finite, we need to discretize the input parameters for the calibration. The laser indices are already discrete, so that we only have to divide the possible values for the ranges and the incidence angles into cells. We decided to divide the incidence angles uniformly, since they are in a fixed interval $[0, \pi/2)$ and the ranges logarithmically, to better be able to cope with different scanner types that have very different maximum ranges. To enforce a smoothness criterion for the calibration function, we add artificial edges between neighboring nodes (one step apart in range or incidence angle). These edges employ the difference of the two values as the error function, thereby requiring them to be similar. One can imagine this as a regular grid, connecting all the nodes in the graph. The stiffness of the grid is determined by the weight of the edges. The higher the weights, the smoother the calibration function has to be. This method also ensures that the complete graph is connected, which is a necessity for the optimization procedure.

In addition to the vertices representing calibration values, there are vertices describing the map of the environment. The map consists of a 3D grid with a fixed cell size and every occupied cell has one corresponding vertex in the graph. Each vertex represents the surface reflecting the laser and we estimate its intensity. The edges represent measurements from the provided point clouds, which are represented as constraints in the graph. Such a constraint connects the corresponding calibration values with the map cell into which the measurement fell. Fig. 3 visualizes the graph structure for a single constraint for the three different modes described in Sec. III-B.

In our concrete implementation we perform two steps that are supposed to reduce the graph size and the noise in the edges in advance. First, we average over the measurements per laser that fall into the same map cell and were observed from a similar perspective and add them as a single combined measurement. Furthermore, we reject map cells where individual laser lines already show a high variance in the remission values, which implies that the surface represented by this cell is not uniformly colored.

E. Expectation Maximization

One problem we have to approach is that the optimization system described above has one trivial solution. If both the calibration parameters and all cell values are zero, the error will also be zero. To circumvent the optimization process from converging to this solution, we chose to perform the optimization in an Expectation-Maximization (EM) fashion and alternate between computing M with fixed C in the expectation step and optimize C with fixed M in the maximization step. Levinson *et al.* [5] used a similar formulation to circumvent the problem of converging to the trivial solution.

Since the calibration parameters can be scaled arbitrarily, we normalize them in each EM-step so that the average calibrated intensity over the whole map is 1.0.

F. Initial Guess

As the initial guess for the calibration parameters we choose only a constant factor for each individual laser, such that the average calibrated remission per laser is 1.0 on the respective dataset.

IV. EXPERIMENTS

We performed experiments on artificial and real data to analyze the properties of our calibration system. We will first present the datasets we used, followed by the experimental evaluation. For all experiments we used 0.5 m for the cell size of the map and divided the incidence angle into 10 cells (9° per cell) and the range into 60 cells, whereas the range of cell i is $(1.1^i - 1)$ meters.

A. Datasets

For our experiments, we used two different robot platforms. Our robot called *Obelix* is shown on the left side of Fig. 4. It is equipped with a multitude of laser scanners, including a downwards facing SICK LMS 151 that accumulates 3D data while driving, a tilting Hokuyo UTM and a Velodyne HDL-32E (a multibeam 3D laser scanner with 32 individual lasers). The second robot is called *Viona* (see Fig. 4 right) and it is equipped with a Velodyne HDL-64E (a multibeam 3D laser scanner with 64 individual lasers). All scanners, which we considered in our experiments, work with 905 nm infrared light.

We captured an outdoor dataset with each robot. Fig. 5 shows the trajectories on aerial images. A good dataset for calibration purposes contains measurements of the same area from different perspectives, meaning from different viewing angles and ranges. This can, e.g., be achieved by driving a robot towards a uniformly colored wall from a far distance, then driving along the wall within a short distance.

From the *Obelix*-dataset we extracted about 680,000 and from the *Viona*-dataset about 500,000 constraints for the optimization procedure.

We used the graph-based SLAM system described in our earlier work [12] to create a globally consistent trajectory to register the point clouds to each other.



Fig. 4. Left: The robot “Obelix” is equipped with three laser scanner types from different manufacturers, with altogether 34 individual lasers. Right: The robot “Viona” is equipped with a Velodyne scanner containing 64 individual lasers.



Fig. 5. The trajectories taken by the robots to acquire the datasets, overlaid on aerial images (©Google). With the robot *Obelix* on the left and with the robot *Viona* on the right.

B. Artificial Remission Values as Ground Truth

To show that our system is able to recover correct calibration parameters given a dataset, we first performed an experiment with artificial remission data on a real 3D dataset. We used the Velodyne points from the *Obelix* dataset and replaced the remission values of all 32 beams (numbered 0-31) by the following arbitrarily chosen function:

$$\begin{aligned} C(l, z_k^l) &= c(r_k^l, \alpha_k^l, l) \cdot i_k^l & (9) \\ &= \frac{0.1 \cdot (l + 1)}{(0.1 \cdot r_k^l + 1) \cdot (\alpha_k^l + 1)^2} \cdot i_k^l, \end{aligned}$$

where we perturb $i_k^l = \mathcal{N}(1, \sigma^2)$ with Gaussian noise ($\sigma = 0.05$).

Fig. 6 shows the results of our calibration procedure in comparison to the ground truth. Please note that the calibration result can always differ by a constant scaling factor on the remission values. Hence, we always normalize the calibration to return an average remission of 1.0 on the learning dataset. The plots in Fig. 6 are appropriately equalized to not differ by such a factor. One can see that our method was able to accurately recover the calibration function with a remaining median error on the dataset of 0.009.

C. Real Datasets

As second experiment, we applied our method on the two datasets described above. We evaluate the three different modes for the optimization as described in Sec. III-B with different independence assumptions. The error distributions before and after optimization for both datasets can be seen in Fig. 7.

One can see that the optimization procedure substantially improves the error compared to the initial guess. For an

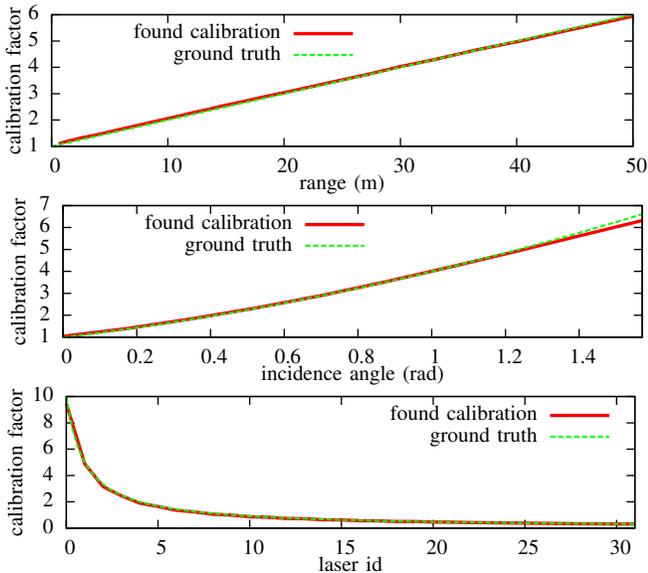


Fig. 6. These plots visualize the found calibration parameters and the ground truth curves for the simulated dataset. The ground truth values are the reciprocal of the corresponding part in Eq. 9, since the purpose of these factors is to revert the influence.

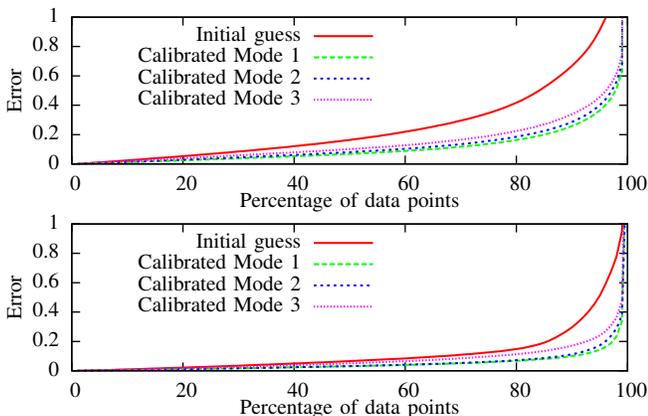


Fig. 7. This plot visualizes the error distribution before and after calibration for the different calibration modes on the Obelix dataset (top) and the Viona dataset (bottom).

average intensity of 1.0, the initial guess provides a median error for the Obelix dataset of 0.164, whereas the median error after optimization is 0.069 (mode 1), 0.081 (mode 2), and 0.101 (mode 3). For the Viona dataset, the initial error is 0.066 and after optimization decreases to 0.031 (mode 1), 0.032 (mode 2), and 0.05 (mode 3). One aspect we can observe in Fig. 7 is that the error increases drastically in the last 20%. This can mainly be attributed to non-uniformly colored, noisy cells in the map, as visible in Fig. 8, where we show the calibrated map for the Obelix dataset and visualize the remaining error. The high remaining error on the lawn does not necessarily imply a bad calibration, just that the assumption for our error metric (uniformly colored flat surface) is not properly fulfilled. The middle image of Fig. 8, as well as the example images in Fig. 1, show the final calibration result for the Obelix dataset. One can clearly distinguish vegetation from concrete and see road markings

in the obtained data. The much smaller final error for the Viona dataset is due to a much friendlier environment for the dataset, with more planar, uniformly colored structures. Besides the expected noise in the measured remission values and the above-mentioned structural issues of real world environments there are additional sources for the remaining errors. These are among other: extreme values for the measured remission (values close to zero and close to the maximum remission are unreliable), the discretization of the returned remission values from the lasers (e.g., only 256 different values for the Velodyne), errors in the registration of the point clouds, and errors in the normal extraction step.

One aspect we can learn from the results in Fig. 7 is that the error difference between mode 1 and mode 2 is minor and can probably be attributed to overfitting in mode 1, meaning that the underlying independence assumption seems to be justified. This means that the influence of the range onto the measured remission is independent to the influence of the incidence angle of the measurement and can therefore be estimated individually. Mode 3 returned the worst result, as expected. While it is still pretty close to the other optimization setups, there seem to be slight differences in the behaviors of the individual laser sensors, even if they are of the same type. Yet, since the risk for overfitting and the needed number of datapoints decreases substantially from mode 1 to mode 3, we decided to use the result of mode 2, whereas we use the result of mode 3 as the initial guess for the optimization procedure.

Fig. 9 shows the found calibration parameters for the laser scanners in our experiments. One can see that most of the scanners behave as expected, with increasing correction factors for increasing ranges and incidence angles. Yet, the Velodyne HDL-32E shows mostly constant factors for different ranges. We believe that the sensor already normalizes the remissions with respect to the measured range. This shows one of the advantages of our methods compared to a physical model for the behavior of the sensor, since we do not need any knowledge of how the measurements are done and what kind of normalizations are already done internally on the sensor. The SICK LMS 151 shows very little influence of the incidence angle. This might be related to the strength of the used laser or an internal calibration procedure. Yet, most likely our dataset did not contain enough evidence for the different angles, since the scanner is mounted in a fixed position on the robot. Therefore the data might have been ambiguous regarding the distribution of the error onto the range influence or the angles influence. The bottom plot shows the individual calibration factors per laser beam. Please note the logarithmic scale in this plot. Here one can see that the uncalibrated remission values have very different value margins.

It remains to show that the learned calibration parameters are not only applicable to the dataset we learned them on. Therefore we captured another dataset with the Obelix robot on our campus in an area that does not overlap with the original dataset, learned calibration parameters on it with mode 2 and applied the learned calibrations crosswise onto

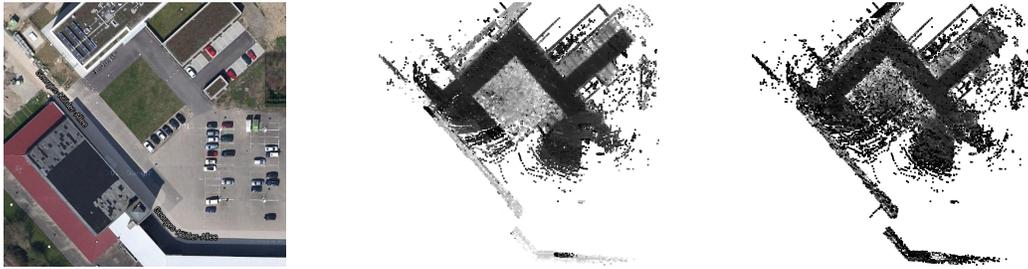


Fig. 8. Aerial image (©Google) of the area where we collected the Obelix dataset (left), the calibrated remission map (middle) and the remaining variances of the measurements per cell, whereas darker values represent lower errors (right). Flat, uniformly colored areas like the street or building walls have very low remaining errors, while the lawn, e.g., still has a high variance, resulting from non-uniform colors and noisy 3D structure.

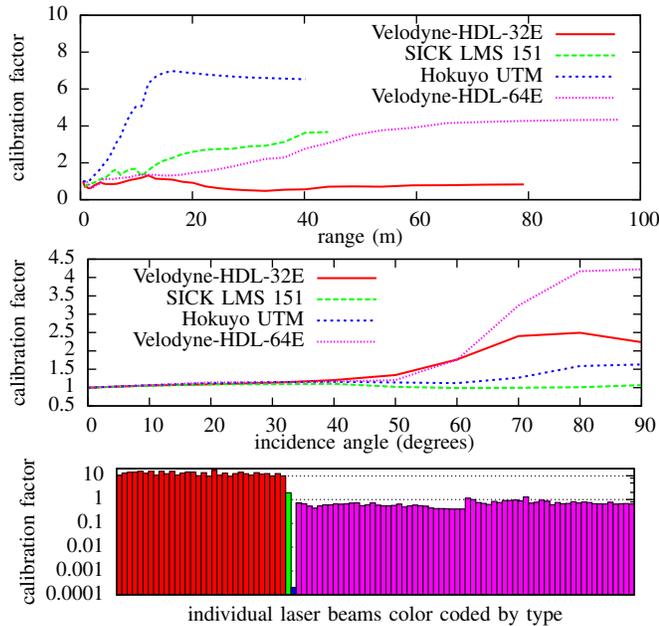


Fig. 9. These plots visualize the found calibration parameters for the two datasets we used. Even though mode 3 returned the worst result, it is better to visualize. In comparison, showing 98 individual plots (one for each individual laser beam) would be impractical. The top plot shows the calibration factors per range for each laser type, the middle plot shows the calibration factors per incidence angle for each laser type and the bottom plot shows the calibration factor for each individual laser, with the same colors as in the plots above. Please note the logarithmic scale on the third plot.

the datasets. The resulting error distributions are shown in Fig. 10. It can be seen that the errors are similar, and mainly differ in the high-error areas, meaning the outliers. This implies that the learned calibration parameters can be generalized to other datasets.

In all our experiments it took less than 15 minutes to get the final calibration parameters given the point clouds on the hard drive, using one thread of a standard Intel I7 PC.

V. CONCLUSIONS

In this paper, we presented an approach for the unsupervised calibration of remission values of a set of heterogeneous laser scanners, based on a dataset of registered point clouds. Our system minimizes the squared error for measurements falling into the same 3D cell to best correspond on the reflectance value of the cell even for different lasers, ranges and incidence angles. We presented experiments with

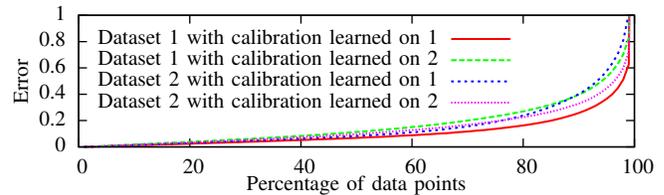


Fig. 10. Error distributions of calibrations learned on two different datasets and applied crosswise to each other.

multiple and different types of laser scanners, showing that after calibration with our system, the lasers can all be used together to create consistent remission maps of the environment.

REFERENCES

- [1] H. Andreasson and A. J. Lilienthal, "Vision aided 3D laser based registration," in *Proceedings of the European Conference on Mobile Robots (ECMR)*, 2007, pp. 192–197.
- [2] S. Kaasalainen, E. Ahokas, J. Hyypä, and J. Suomalainen, "Study of surface brightness from backscattered laser intensity: calibration of laser data," *Geoscience and Remote Sensing Letters, IEEE*, vol. 2, no. 3, pp. 255–259, 2005.
- [3] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard, "g2o: A general framework for graph optimization," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2011.
- [4] J. Levinson and S. Thrun, "Robust vehicle localization in urban environments using probabilistic maps," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [5] —, "Unsupervised calibration for multi-beam lasers," in *Proceedings of the International Symposium on Experimental Robotics (ISER)*, 2010, pp. 179–193.
- [6] W. Maddern, A. Harrison, and P. Newman, "Lost in translation (and rotation): Fast extrinsic calibration for 2D and 3D LIDARs," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2012.
- [7] C. McManus, P. Furgale, B. Stenning, and T. D. Barfoot, "Lighting-invariant visual teach and repeat using appearance-based lidar," *Journal of Field Robotics*, vol. 30, no. 2, pp. 254–287, 2013.
- [8] N. Muhammad and S. Lacroix, "Calibration of a rotating multi-beam lidar," in *Proc. of the Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010.
- [9] M. Sheehan, A. Harrison, and P. Newman, "Self-calibration for a 3D laser," *Int. Journal of Robotics Research*, vol. 21, no. 5, 2012.
- [10] S. Soudarissanane, R. Lindenbergh, M. Menenti, and P. Teunissen, "Incidence angle influence on the quality of terrestrial laser scanning points," in *Proc. of the ISPRS Workshop Laserscanning*, 2009.
- [11] "Open source software released by Bastian Steder," www.informatik.uni-freiburg.de/~steder/openSourceSoftware.html.
- [12] B. Steder, M. Ruhnke, S. Grzonka, and W. Burgard, "Place recognition in 3D scans using a combination of bag of words and point feature based relative pose estimation," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011.
- [13] K. M. Wurm, H. Kretschmar, R. Kümmerle, C. Stachniss, and W. Burgard, "Identifying vegetation from laser data in structured outdoor environments," *Journal of Robotics & Autonomous Systems*, vol. 62, no. 5, pp. 675–684, 2014.