

Automatic Extrinsic Calibration of Multiple Laser Range Sensors with Little Overlap

Jörg Röwekämper Michael Ruhnke Bastian Steder Wolfram Burgard Gian Diego Tipaldi

Abstract—Networks of laser range finders are a popular tool for monitoring large cluttered areas and to track people. Whenever multiple scanners are used for this purpose, one major problem is how to determine the relative positions of all the scanners. In this paper, we present a novel approach to calibrate a network of multiple planar laser range finders scanning horizontally. To robustly deal with the potentially restricted overlap between the fields of view, our approach only requires a dynamic object, e.g., a person, moving through the observed area. We employ a RANSAC-like algorithm to find the correspondences between the measurements of the different laser range finders. Based on these correspondences we formulate a graph-based optimization problem to determine the maximum likelihood extrinsic parameters of the sensor network. Furthermore, we present a method to evaluate the consistency of the resulting calibration based on visibility constraints. Experiments on real and simulated data show that the proposed approach yields better results than techniques that only perform pairwise calibration.

I. INTRODUCTION

Laser range finders are widely used sensors in robotics and in industry. Applications often aim to acquire spatial information about the environment, to avoid obstacles, to build maps or to localize a mobile robot in a map. They have been further used to track people, to monitor areas or to implement safety zones, i.e., to detect humans and switch off dangerous machines when they are too close. For surveying a larger area, typically a network of multiple laser range finders is necessary. At the same time, for financial reasons, one seeks to minimize the number of sensors in the network, thereby reducing the overlap between their field of view.

Calibrating such a sensor network can be tedious, especially when the overlap between the visible areas of the individual scanners is small. Standard calibration algorithms often rely on predefined patterns or calibration aids [11, 22, 16]. These calibration patterns solve the perceptual aliasing problem and improve the data association between the observations. Afterwards, these methods employ scan matching algorithms to align the associated measurements and estimate the extrinsic parameters. In practice, the use of dedicated calibration aids is not helpful in the context of laser range scanners. At the same time, without aids, data association is more problematic which directly influences the calibration results as scan matching algorithms require a good initial guess.

All authors are with the Department of Computer Science at the University of Freiburg, Germany. This work has partly been supported by the European Commission under FP7-610917-STAMINA.




Fig. 1. Top: Experimental setup with five laser range finders with unknown poses. Bottom: Calibrated network using our approach. The colored wavefronts show the field of view of the individual laser range finders (depicted as dots of the same color). The size of the environment is $27 \times 11.5 m^2$.

In this paper we present a novel approach to calibrate a network of multiple stationary, horizontally scanning 2D laser range finders with limited overlap. We formulate the extrinsic calibration problem as a joint maximum likelihood estimation problem, in which we jointly estimate the extrinsic parameters of all the laser range finders at once. We model the optimization in a similar fashion to bundle adjustment [21] and seek to minimize the reprojection error of corresponding measurements from all laser range finder pairs and all timestamps. To address the problem of gauge freedom and non-observability of the whole system, we fix one range finder to stay at the center of the global reference frame. Our approach differs from most other calibration techniques, which perform only pairwise calibration between one master sensor and all the others. We do not rely on any predefined calibration pattern and we automatically compute a valid initial guess for the nonlinear optimization problem exploiting the moving object. For providing a valid initial guess to the nonlinear least square problem, we assume that a moving object, e.g., a person, is present in the environment. Furthermore, we developed a consistency measure for the final calibration result, which we use to analyze the calibration results and improve them. We furthermore assume that the graph built from the overlaps between the visibility areas of the individual scanners is connected, and that a

single dynamic object is moving through all overlapping areas, i.e., traversing every edge of the graph. We do not assume a specific laser range finder, nor we require a specific resolution, frequency or field of view. The sensor network can contain different types of laser range finders and only requires that they are mounted horizontally and that they provide consistent timestamps for their measurements.

II. RELATED WORK

Sensor calibration is a well studied field and many approaches have been developed in the last decades. In this section, we mainly focus on approaches that perform calibration of multiple sensors without the need of a calibration pattern. We will not cover calibration algorithms for only intrinsic parameters, nor calibration of an individual sensor relative to the coordinate frame of a robot.

Many researchers focused on calibrating cameras with respect to the laser range finder. For example, Zhang and Pless [22] proposed a tool for camera-laser calibration. They use a planar checkerboard pattern and to seek for correspondences between the observations of a camera and a laser range finder. They propose a direct solution and a non-linear refinement step. Pandey et al. [16] extended the work to 3D using a panoramic camera and a velodyne scanner. Zhao et al. [23] presented a pairwise calibration system between multiple laser range finders and cameras. They perform the calibration of the sensors with respect to a common reference frame on the car and rely on a given pairwise data association between points in the camera and points in the scan. Mirzaei et al. [15] propose an approach to jointly estimate the intrinsic parameters of a 3D laser range finder and its extrinsic parameters with respect to a camera. They divide the problem into two separate least square minimization tasks and find a closed-form solution to precisely determine an initial guess. Levinson and Thrun [12] exploit depth discontinuities in laser data and edges in images to estimate the 6-DoF transformation between a Velodyne sensor and a camera. All these methods assume that the involved sensors cover the same area and thus have largely overlapping fields of view.

Other work focused on calibrating multiple laser range finders with respect to themselves in a pairwise fashion. Choi et al. [6] calibrate a pair of 2D laser range finders observing different, non co-planar planes. Their approach relies on a calibration structure consisting of two orthogonal planes and uses a non-linear least squares formulation to compute the calibration parameters. Our approach does not need the specific calibration structure but we assume that the sensors are mounted in one plane. Sheehan et al. [19] calibrate a set of 2D laser range finders mounted on a spinning disk. They employ the Renyi entropy as a measure of point cloud quality and cast the problem as entropy minimization. The same idea has later been used to calibrate the resulting 3D spinning laser range finder with an inertial measurement unit and a set of 2D laser range finders [13]. The authors also provide an approximation that gracefully degrades with the amount of computation used. Their approach assumes that the robot

moves and while moving all sensors will observe the same part of the environment. Le and Ng [11] jointly calibrate intrinsic and extrinsic parameters of generalized 3D sensors, e.g., a 3D laser range finder, stereo cameras or a rotating 2D laser range finders. The approach relies on matching point correspondences between pairwise sensors and computes the extrinsic parameters via least square optimization. They calibrate networks of sensors sequentially by chaining their respective transformations. A main limitation is the requirement of a checkerboard for point correspondences.

In contrast to these approaches, we address calibration as a global optimization problem and do not rely on the pairwise calibration of sensors. Similar ideas have been proposed recently by other researchers. Brookshire and Teller [3] recover the 2D rigid transformation between pairs of sensors mounted on the same plane. They only require incremental pose measurements and can use asynchronous sensors (but with time-sync between them). They also compute the Cramer-Rao lower bound of the calibration and estimate the uncertainty of the relative transformation. The authors use non-linear least square optimization and reach a submillimeter accuracy. The authors recently extended their method to 3D [4] by performing the minimization and the interpolation using dual-quaternions and Lie-algebra. Schneider et al. [18] present a calibration algorithm based on the sensor odometry. Given the time-synchronized delta poses of two sensors, they estimate the relative pose between them using the unscented Kalman filter. Kümmerle et al. [9] present a similar approach, also based on the odometry estimation of different sensors. They propose to simultaneously estimate the parameters of the sensors simultaneously to determining the position of the robot and the map of the environment. The main difference between their and our approach is that we consider stationary sensors and, hence, we cannot rely on computing incremental pose measurements. The limited overlap results in a reduction of observable data by the sensor, therefore it is difficult to reach the same level of accuracy as of the previous approaches.

Recently, Fernández-Moral et al. [7] presented a calibration approach based on matching plane observations from different range sensors. Their method poses no restrictions on the position of the cameras and only assumes that there is a planar surface observed simultaneously. They formulate the calibration jointly between all the sensors, using a non-linear least square formulation. Although their approach shares similarities to ours, we do not require the presence of planar surfaces in the environment and we also provide a consistency measure of the resulting calibration.

III. ESTIMATING A VALID INITIAL GUESS

In this section we will describe how we compute the initial guess for the least square problem. We proceed in two steps. First, we compute a guess between pairs of laser range finders that have an overlapping region. Second, we transform the estimated relative transformations of the first step to an optimization problem to compute a joint initial guess for the whole system of scanners.

A. Computation of the Pairwise Initial Guess

To figure out which laser range finders jointly observe a part of the environment, we exploit the presence of dynamic obstacles moving through the environment. In the initialization phase, we estimate, for each laser range finder, the static part of the environment using background subtraction techniques [2]. When the moving object is entering the field of view of one sensor, we exploit the learned background to identify the measurements belonging to the moving person. Based on the dynamic range measurements introduced by the moving object, we compute the pairwise relative transformation between corresponding laser range finders with RANSAC [8]. In detail, for each laser range finder, we collect the set of measurements belonging to the moving object and store them together with their time stamps. Let $\mathcal{P}_i = \{\mathbf{p}_{i,q}^t\}$ be the set of measured points of the laser range finder i where $\mathbf{p}_{i,q}^t$ is the q -th point recorded at time t . Let $\mathcal{T}_{i,j} = \{t\}$ be the set of time stamps in which the laser range finder i and j have observed the moving object at the same time t . For each pair of laser range finders i and j , we sample the minimal hypothesis set in a two stage process. At first, we sample a time instance t from the set $\mathcal{T}_{i,j}$. Then we sample two random points from each set \mathcal{P}_i and \mathcal{P}_j recorded at time t . After calculating an initial transformation from these two correspondences, we compute the number of inliers over all time instances. We consider a point an inlier if the distance between the point and its nearest neighbor in the other scan is below a specified threshold according to the computed transformation. We repeat this process multiple times and in the end return the transformation hypothesis with the highest inlier count as initial guess.

To avoid false initializations, it is of utmost importance to use only observations from the dynamic object. The overlapping area can be rather small, so that a large fraction of the point correspondences from the static scene will be outliers. In this case, RANSAC needs many more iterations to converge and the number of outliers could surpass its breaking point, which may lead to a biased result. One might argue that alternative and more sophisticated matching algorithms might overcome this problem [20, 17]. However, we found out that relying on just the dynamic object measurements suffices in most of the cases.

B. Computation of the Joint Initial Guess

The parameters estimated in the previous section only relates pairs of laser range finders relative to each other. In order to properly initialize the optimization problem, we need to combine these transformations. A naive approach would be to compute the graph from the overlapping areas and chain the transformations along its minimum spanning tree. However, the errors typically accumulate while chaining the transformations, potentially resulting in poor point correspondences between the measured points. To solve

¹In practice, there will be always some differences in the time stamps. In this paper we consider two times tamps equal if they differ by less than half of the period between to scans. Which is 10 ms for a frequency of 50 Hz.




Fig. 2. Example factor graph using pose to pose measurements. We measure the relative poses $\mathbf{z}_{ij} = \mathbf{x}_i \ominus \mathbf{x}_j$ between four laser range finders using RANSAC (white nodes). The factors correspond to the Measurements (black squares).




Fig. 3. Example factor graph using point landmarks. The positions of the landmarks \mathbf{l}_k (shaded nodes) implicitly provide measurements \mathbf{z}_{ij} (black squares) between the laser range finders \mathbf{x}_i (white nodes).

this problem, we seek to find the maximum likelihood configuration of all sensors. To achieve this, we construct a factor graph (see Fig. 2) where each node represents a laser range finder and each factor represents the transformation estimated with RANSAC. We scale the transformation uncertainties by the inverse of the number of inliers computed by RANSAC. Intuitively, this is equivalent to believe more in transformations that have more support from the inlier set. More formally, let $\mathbf{x}_i, \mathbf{x}_j, \mathbf{z}_{i,j} \in SE(2)$ be the pose of the i -th and j -th range finder, and their transformation estimated by RANSAC, respectively. We solve for the initial global alignment by minimizing the following least squares problem:

$$(\mathbf{x}_1, \dots, \mathbf{x}_N)^* = \operatorname{argmin}_{\mathbf{x}_1, \dots, \mathbf{x}_N} \sum_{\mathbf{z}_{i,j}} \|\mathbf{x}_j \ominus \mathbf{x}_i - \mathbf{z}_{i,j}\|_{\Omega_{i,j}}^2 \quad (1)$$

where $\Omega_{i,j}$ denotes the information matrix associated with the measurement $\mathbf{z}_{i,j}$. The above formulation is closely related to the simultaneous localization and mapping (SLAM) problem, formulated as pose graph optimization. Although the two problems share the same error functions for the optimization, they differ in the type of data used to compute the pairwise transformations. In a SLAM setting, there is one robot going around the environment collecting data, while in our case the sensor positions do not change. At the same time, the sensors have only a limited overlap. In many cases, the extrinsic parameters computed by this approach are not accurate enough for the task at hand. The transformation returned by RANSAC are often computed using only a small subset of the points, due to noise in sensor data, non synchronized sensors and slightly different viewpoints. Therefore, we rely on this result as a good initial guess for the main calibration procedure described in the following section.

IV. JOINT CALIBRATION OF MULTIPLE LASER RANGE FINDERS

In this section, we describe the proposed calibration scheme and we show how we can compute a measure of consistency of the final calibration. We estimate the relative transformations of all laser range finder pairs by constructing a global optimization problem in which we include all point correspondences for all laser pairs and all timestamps. In correspondence to the SLAM problem, each point correspondence can be regarded as a landmark, observed from two different laser range finders at one specific time.

We build another factor graph, similar to the one described in the previous section. This one, however, includes additional nodes to model the point correspondences observed by the laser range finders. Accordingly, each factor expresses the spatial relationships between two sensor poses and a jointly observed landmark. Figure 3 depicts the structure of this graph. Note the difference to the model in Figure 2. As can be seen, each pair of poses in the new graph can have multiple edges between each other, see for instance the nodes \mathbf{x}_3 and \mathbf{x}_4 . Note that these landmarks are neither distinctive features nor come from static observations since we observe only points sampled from an moving object. Therefore, there is no re-observation of landmarks in a second time instance.

Given an initial guess, we build the factor graph as following. For each pair of sensors i and j , we express the measured points in the sets \mathcal{P}_i and \mathcal{P}_j in the local coordinate frames of their corresponding sensor. For each point we compute a local surface approximation and store the corresponding normal direction. Then, for each time stamp t , we associate the points $\mathbf{p}_{i,q}^t$ and $\mathbf{p}_{j,r}^t$ if their euclidean and angular distances are below a threshold. For each associated point we insert a landmark node \mathbf{l}_l in the graph and its associated factor $\mathbf{z}_{i,j,l}$.

For each point correspondence $\langle \mathbf{p}_{i,q}^t, \mathbf{p}_{j,r}^t \rangle$ we look for the configuration of poses \mathbf{x}_i and \mathbf{x}_j that minimizes the reprojection error

$$\mathbf{z}_{i,j,l} = \mathbf{p}_{i,q} - (\mathbf{x}_j \ominus \mathbf{x}_i) \oplus \mathbf{p}_{j,r}, \quad (2)$$

where q and r are the point indices associated with the l -th landmark (note that we omitted the time index for notational simplicity). To compute the maximum-likelihood calibration parameters $(\mathbf{x}_1, \dots, \mathbf{x}_N)^*$, we solve the least squares problem

$$(\mathbf{x}_1, \dots, \mathbf{x}_N)^* = \operatorname{argmin}_{\mathbf{x}_1, \dots, \mathbf{x}_N} \sum_{\mathbf{z}_{i,j,l}} \|\mathbf{z}_{i,j,l}\|_{\Omega_{i,j,l}}^2, \quad (3)$$

where $\Omega_{i,j,l}$ denotes the information matrix associated with the measurement $\mathbf{z}_{i,j,l}$.

We iteratively repeat this process and decrease the distance thresholds for point correspondences with each iteration. After the optimization problem converged, we perform a final refinement step in which we also include the observations from the static scene. Note that there is no guarantee that the proposed optimization procedure finds the global minimum for all possible initializations. Still, we found out in our




Fig. 4. Simulation environment with seventeen laser range finders. Each a wave front indicates the field of view of the corresponding laser range finder. Each one starts at the origin of the corresponding laser range finder indicated with a red dot. Note that the overlaps of the fields of view of some of the laser range finders are rather small which increases the difficulty of estimating the correct transformations.

experiments that the proposed method typically converges to the global minimum given the jointly computed initial guess, even in real world settings.

A. Consistency and Uncertainty of the Resulting Calibration

An important aspect of calibration algorithm is to provide the user with a measure of how reliable the resulting calibration is. Most approaches provide observability analysis and estimate the Cramer Rao lower bound of the system. However, as also pointed out by Censi [5], computing such a bound is not trivial, not even in the simple case of only two sensors. Moreover, there might be some unobservable bias due to wrong point correspondences.

In light of this aspects, we decided to focus more on understanding whether the final estimate corresponds to a local minimum. To achieve this, we extend the work of Mazuran et al. [14] to handle observations taken at different times from the same pose. We achieve this by including an exterior loop in the algorithm, where we compute the consistency measure on each time stamp. Finally, we compute for each pair of laser range finders a consistency value by projecting the measurement of one sensor into the coordinate frame of the other and computing the ratio between the number of consistent observations vs. the number of total observations. This provides an more intuitive measure to non expert users. We then visualize the transformations with lowest consistency value, together with the estimated overlapping area. We use this information to guide the process and collect more data on the low consistency region.

V. EXPERIMENTS

We evaluated the proposed calibration method on both simulated and real datasets. For the simulated data, we considered 17 laser range finders in an environment of $20 \times 10 m^2$ and simulated a person walking through it. Figure 4 shows the corresponding environment and the visible area of each laser range finder. Wavefronts indicate the field of view of each sensor. The simulated laser range finders have a resolution of 0.5° , a field of view of 190° , a max range of $15 m$ and a Gaussian noise with $0.015 m$ standard deviation.

For the real world data, we used five laser range finders in a reconfigurable experimental area and recorded three data

sets with a person walking along different trajectories. To estimate the ground truth, we first removed all the walls in the environment to have full overlap between the laser range finders. Then, we estimated their relative transformations using ICP over multiple scans and averaged the resulting estimate. After estimating their transformation, we put the walls back in their position. The bottom of Figure 1 shows the experimental setup after we inserted the walls. As in the previous figures, we use colored wavefronts to indicate the field of view of each sensor.

The lower right pair of laser range finders in Figure 1 is mounted on different heights therefore the laser range finders see different parts of the static environment. This makes it more difficult to estimate the transformation because the laser range finders have nearly no overlap in the static parts of the scans. Also not perfectly adjusted tilt angles and time stamping errors of the laser range scans have an influence on the accuracy on the experiment results. The simulated data has the advantage that we know the true pose of the laser range finders and that we can avoid biases due to tilted laser range finders and inaccurate time stamping of scans.

We compared the accuracy of the proposed calibration with respect to the RANSAC initial guess (RANSAC) and a pairwise iterative closest point (ICP) algorithm. We performed the pairwise ICP only with the static scans and use the RANSAC results as an initial guess. In the error metric, we consider the pairwise transformations instead of the global pose of the laser range finders, which could be influenced by the choice of the reference frame. Because of the non-deterministic nature of RANSAC we ran all experiment 100 times and computed mean, standard deviation and maximum values of the error metric. We performed all experiments with the same set of parameters. For RANSAC, we used an inlier distance of 0.05 m , a confidence level of 0.23 , and an inlier probability of 0.25 . During the ICP iterations we decreased the correspondence search radius from 0.7 m down to 0.1 m . For both optimization problems, we used the Levenberg-Marquardt algorithm implemented in g^2o [10] with dynamic covariance scaling [1] as robust kernel. We used thresholds ranging from 0.7 m to 0.1 m in 7 iterations for the point correspondences and 0.05 m for the final refinement using static measurements.

A. Results

Figure 5 show the results of the experiments. The simulated runs have the labels Set 1, Set 2, and Set 3 and the real world runs have the labels Set 4, Set 5, and Set 6. The shortest simulated data set spans 299 seconds and the longest 432 second. The shortest real world data sets corresponds to 58 seconds and the longest to 84 seconds.

For the simulation sets, RANSAC followed by global optimization achieves a mean translational error of about 0.055 m . The proposed approach, when using only dynamic data, achieves an accuracy of 0.025 m , which improves to 0.014 m when we also consider the static structure. The standard deviation for our approach is 0.0002 m , indicating that all runs converge to the same minimum. For the rotational

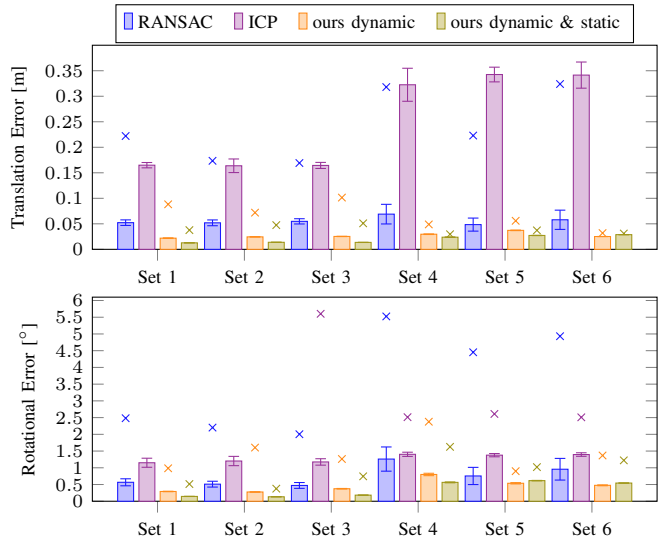


Fig. 5. Errors of the simulated (Set 1, Set 2, Set 3) and experimental (Set 4, Set 5, Set 6) data sets. The boxes indicate the mean error of the averaged results over all runs and the bars show the standard deviation to illustrate the distribution of the error over the runs. The crosses represent the maximum errors of all relative transformation over all runs. The top figure shows the translational errors and the bottom figure the rotational errors. All ICP max translational errors are above 0.48 m for the simulated sets and above 0.86 m for the experimental sets. The ICP max rotational errors for Set 1 and Set 2 are 17.2° and 13.8°

errors, RANSAC followed by optimization achieves a mean error of about 0.5° . Our approach using dynamic data has a similar accuracy of 0.37° . This improves drastically when including the static structure to 0.18° . We believe this is due to the fact that the dynamic data is often concentrated in a small portion of the environment, poorly constraining the angular part of the transformation. ICP performs worse than RANSAC and our approach. We believe this is caused by the limited number of point correspondences from static objects, which are not enough to compensate for the noise.

The results obtained with the real data show a similar trend. The main difference is that the improvement in accuracy when including the static structure is not as drastic as before. We believe this is due to the imperfect planar alignment of the sensors and because the sensors are mounted at different heights. When we closely inspect the bottom right part of Figure 1, we see that one laser range finder is measuring a table (the orange sensor) while another one is only measuring the wall, indicating that it is mounted above the table. The experimental results demonstrate that our approach improves over a naive approach using RANSAC followed by global optimization on pairwise measurement only. The results also show that using the information from the static scene improves the accuracy, even if there is only limited overlap. To give an idea about the runtime, for the longest simulated set, it took 67.8 s to compute the joint initial guess and 4.2 s to optimize. The set contains the laser range readings of 17 scanner and is 432 s long.

B. Consistency Check

In this section we present the advantage and effectiveness of the proposed consistency measure. To do so, we evaluate

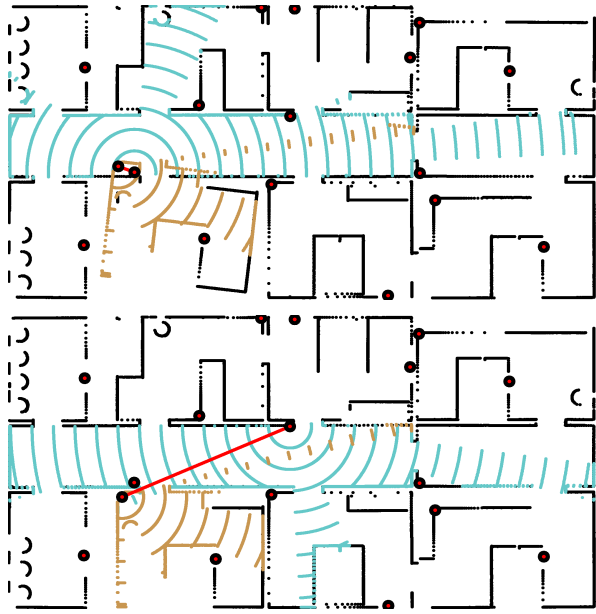


Fig. 6. Example for the application of the consistency measure. The top figure shows the calibrated laser poses, whereas one of the lasers (marked in brown) is not well aligned with the rest. The bottom image shows the correct registration of this sensor. Our consistency measure is 0.24 for the link between the brown and the blue laser in the top image, whereas the same measure in the bottom image is 0.92, thereby correctly capturing the bad alignment in the top image.

our consistency measure on two simulated runs. The first consists of a data set in which we purposely did not cover the whole overlapping area with the moving object. For the second, we use the data collected for Set 2. Figure 6 shows the two different datasets, where we highlight the transformation with the lowest consistency value. In the first run (top), we obtain a consistency score of 0.24, clearly indicating that we cannot trust the resulting transformation. For the second run (bottom), we have instead a consistency score of 0.92, which indicates a consistent calibration. Thus, the end user can use this information to inspect the final calibration and decide if he needs further data to achieve desired target accuracy.

VI. CONCLUSION

In this paper, we presented a novel calibration algorithm for the extrinsic parameters of a set of static laser range finders placed at the same height and scanning horizontally. In contrast to previous approaches, we do not consider pairwise calibration only but rather consider it as a joint maximum likelihood estimation problem. Our approach does not need any calibration pattern and requires only a single object moving through the environment. We furthermore propose an initialization scheme to provide a valid guess and a measure of consistency to automatically check if the calibration is correct. We implemented our approach and evaluated it in both real and simulated data. The results show that our algorithm is able to estimate the extrinsic parameters of the network with high accuracy, even when the fields of view of the individual scanners have very little overlap. Additionally,

our approach can achieve higher accuracy than pairwise calibration, even after a global minimization step. Finally, we show that the proposed consistency measure is able to capture wrong calibrations and is a useful tool for the user.

REFERENCES

- [1] P. Agarwal, G. D. Tipaldi, L. Spinello, C. Stachniss, and W. Burgard. Robust map optimization using dynamic covariance scaling. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2013.
- [2] Y. Benezeth, P.-M. Jodoin, B. Emile, H. Laurent, and C. Rosenberger. Review and evaluation of commonly-implemented background subtraction algorithms. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference*, 2008.
- [3] J. Brookshire and S. Teller. Automatic calibration of multiple coplanar sensors. In *Proc. of Robotics: Science and Systems (RSS)*, 2011.
- [4] J. Brookshire and S. Teller. Extrinsic calibration from per-sensor egomotion. In *Proc. of Robotics: Science and Systems (RSS)*, 2012.
- [5] A. Censi. On achievable accuracy for pose tracking. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2009.
- [6] D.-G. Choi, Y. Bok, J.-S. Kim, and I. S. Kweon. Extrinsic calibration of 2d laser sensors. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2014.
- [7] E. Fernández-Moral, J. González-Jiménez, P. Rives, V. Arévalo, et al. Extrinsic calibration of a set of range cameras in 5 seconds without pattern. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
- [8] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Comm. of the ACM*, 24(6):381–395, 1981.
- [9] R. Kümmerle, G. Grisetti, and W. Burgard. Simultaneous parameter calibration, localization, and mapping. *Advanced Robotics*, 26(17):2021–2041, 2012. doi: 10.1080/01691864.2012.728694.
- [10] R. Kümmerle, G. Grisetti, H. Strasdat, K. Konolige, and W. Burgard. g^2o : A general framework for graph optimization. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2011.
- [11] Q. V. Le and A. Y. Ng. Joint calibration of multiple sensors. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2009.
- [12] J. Levinson and S. Thrun. Automatic calibration of cameras and lasers in arbitrary environments. In *International Symposium on Experimental Robotics*, 2012.
- [13] W. Maddern, A. Harrison, and P. Newman. Lost in translation (and rotation): Rapid extrinsic calibration for 2d and 3d lidars. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2012.
- [14] M. Mazuran, G. D. Tipaldi, L. Spinello, W. Burgard, and C. Stachniss. A Statistical Measure for Map Consistency in SLAM. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2014.
- [15] F. M. Mirzaei, D. G. Kottas, and S. I. Roumeliotis. 3d lidar-camera intrinsic and extrinsic calibration: Identifiability and analytical least-squares-based initialization. *Int. Journal of Robotics Research*, 31(4):452–467, 2012.
- [16] G. Pandey, J. McBride, S. Savarese, and R. Eustice. Extrinsic calibration of a 3d laser scanner and an omnidirectional camera. In *IFAC Symp. on Intelligent Autonomous Vehicles*, volume 7, 2010.
- [17] F. T. Ramos, D. Fox, and H. F. Durrant-Whyte. Crf-matching: Conditional random fields for feature-based scan matching. In *Proc. of Robotics: Science and Systems (RSS)*, 2007.
- [18] S. Schneider, T. Luettel, and H.-J. Wuensche. Odometry-based online extrinsic sensor calibration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2013.
- [19] M. Sheehan, A. Harrison, and P. Newman. Self-calibration for a 3d laser. *Int. Journal of Robotics Research*, 31(5):675–687, 2012.
- [20] G. D. Tipaldi and K. O. Arras. FLIRT – Interest Regions for 2D Range Data. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2010.
- [21] B. Triggs, P. F. McLauchlan, R. I. Hartley, and A. W. Fitzgibbon. Bundle adjustment: modern synthesis. In *Vision algorithms: theory and practice*, pages 298–372. Springer, 2000.
- [22] Q. Zhang and R. Pless. Extrinsic calibration of a camera and laser range finder (improves camera calibration). In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2004.
- [23] H. Zhao, Y. Chen, and R. Shibasaki. An efficient extrinsic calibration of a multiple laser scanners and cameras’ sensor system on a mobile platform. In *Intelligent Vehicles Symposium, 2007 IEEE*, 2007.