

MAXIMUM CONSENSUS LOCALIZATION USING LIDAR SENSORS

J.Axmann¹, C. Brenner¹

¹ Institute of Cartography and Geoinformatics, Leibniz University Hannover, Germany
- (jeldrik.axmann, claus.brenner)@ikg.uni-hannover.de

Commission II, WG II/1

KEY WORDS: Localization, LiDAR, Maximum Consensus, Robust Estimation, Point Cloud Registration, Integrity.

ABSTRACT:

Real world localization tasks based on LiDAR usually face a high proportion of outliers arising from erroneous measurements and changing environments. However, applications such as autonomous driving require a high integrity in all of their components, including localization. Standard localization approaches are often based on (recursive) least squares estimation, for example, using Kalman filters. Since least squares minimization shows a strong susceptibility to outliers, it is not robust.

In this paper, we focus on high integrity vehicle localization and investigate a maximum consensus localization strategy. For our work, we use 2975 epochs from a Velodyne VLP-16 scanner (representing the vehicle scan data), and map data obtained using a Riegl VMX-250 mobile mapping system. We investigate the effects of varying scene geometry on the maximum consensus result by exhaustively computing the consensus values for the entire search space. We analyze the deviations in position and heading for a circular course in a downtown area by comparing the estimation results to a reference trajectory, and show the robustness of the maximum consensus localization.

1. INTRODUCTION

Tackling localization tasks including the fusion of different sensor measurements using Kalman filters is a common approach (Chen, 2012). Applying Kalman filters, uncertainties are modelled using normal distributions and residuals are minimized in a least squares manner (Kalman, 1960). Consequently, there is a strong susceptibility to outliers, which contradict the underlying assumption of independent and identically normal distributed residuals (Chin and Suter, 2017). Even a small amount of outliers can affect the least squares estimate heavily. In real world localization, especially using images or LiDAR sensors, there is usually a large amount of outliers due to erroneous measurements, false assumptions about the ego position, and changing environments.

Estimation techniques based on the maximum consensus criterion can be used to approach this lack of robustness coming along with least squares. The consensus set size serves as robust objective function since the maximum consensus solution corresponds to the candidate model yielding the largest consensus set, which in turn represents the highest consistency with the postulated functional model. In contrast to least squares, maximum consensus methods exclude data beyond a specific threshold preventing these from influencing the estimate. Therefore, maximum consensus techniques are inherently well-suited for parameter estimation in outlier contaminated data (Chin and Suter, 2017). There are various maximum consensus approaches, which can be categorized into heuristic and exact ones.

Whereas heuristic techniques are often able to generate results close to the optimum in a reasonable amount of time, they cannot give any guarantees. The probably most famous heuristic maximum consensus technique is *Random Sample Consensus* introduced by (Fischler and Bolles, 1981). The algorithm randomly draws samples used to instantiate candidate models and

evaluates those by comparing their consensus set sizes. The basic assumption is that within a particular number of trials, there is at least one in which only inliers are picked. The number of trials is computed from the assumed percentage of outliers and a failure probability set by the user. Even though this method has been applied very successfully in many estimation settings, it does not give any guarantees. In addition, the required number of draws is approximately inversely proportional to the probability of a good draw (of a model), mitigating the advantage of lower computational effort in strongly outlier contaminated data.

If gradients are available, non-randomized methods such as gradient ascent or descent seem attractive, however those cannot be used in the non-convex settings frequently encountered. Two deterministic, heuristic approaches, based on a non-smooth penalization and the alternating direction method of multipliers, respectively, have recently been reported by (Le et al., 2017). While they show good results and attractive execution times, especially at higher outlier rates, they do neither ensure global optimality nor being within certain bounds from the optimal result.

Exact algorithms guarantee to obtain the globally optimal result, but they are computationally costly. Since maximum consensus is fundamentally NP-hard (Chin et al., 2018), the optimization cannot be solved in polynomial time. Exact algorithms aim to keep the computational effort low by conducting search strategies (Chin and Suter, 2017). For example, *Branch-and-Bound* (BnB)-, tree-, and exhaustive search were reported in the literature. (Li, 2009) and (Zheng et al., 2011) as well as (Breuel, 1992), (Chin et al., 2014), and (Parra Bustos et al., 2016) introduced BnB algorithms. (Olsson et al., 2008) and (Enqvist et al., 2012) pursued exhaustive search by enumeration, while (Chin et al., 2017) introduced a tree search strategy. In general, search strategies will succeed in acceptable time only if there are very few local maxima in search space, so that only a very small portion of the total search space needs to be explored.

Considering the applications of approximate maximum consensus approaches, the approaches of (Zheng et al., 2011; Li, 2009; Chin et al., 2017; Olsson et al., 2008; Enqvist et al., 2012) focus on robust fitting of feature points in images.

Localization using LiDAR measurements is closely related to point cloud registration. If the points are known in a sensor coordinate system, then aligning point clouds also yields the relative orientation of the sensor standpoints. Point cloud registration is a very well-researched topic. For example, (Chen et al., 1999) and (Aiger et al., 2008) use randomized heuristics, (Li and Hartley, 2007) and (Yang et al., 2013) are using maximum likelihood frameworks, and (Chin et al., 2014) and (Parra Bustos et al., 2016) are using a geometric matching criterion, striving for globally optimal point cloud registration using BnB. The most common algorithm used for point cloud registration is arguably the *Iterative Closest Point* (ICP) algorithm (Besl and McKay, 1992), which approaches intractability by using two alternating phases, (i) finding correspondences and (ii) minimizing the residuals. If least squares is used for the minimization step, this of course also introduces a sensitivity to outliers, which is usually mitigated by additional selection, matching or weighting criteria in the correspondence or minimization phase (Chetverikov et al., 2002). However, even if the optimum solution could be guaranteed in the minimization phase, the algorithm could still converge to a local optimum due to wrong correspondences. Typically, one hopes for a convergence towards the global optimum if the initial values are ‘good enough’.

In this paper, we investigate high integrity localization based on LiDAR, using maximum consensus. Even though in general, the optimization problem has exponential time complexity in the number of variables, it can of course be tackled by exhaustive search in case of a low-dimensional problem. In addition, exhaustive computation of the search space allows us to gain knowledge about typical and atypical situations, and the influence of the geometry of the scene. In detail, in this paper we show (i) a localization based on the registration of two point clouds, a sparse, ‘car sensor’ cloud from a Velodyne VLP-16 scanner, and a dense, high resolution ‘map’ point cloud from a mobile mapping campaign, (ii) a high integrity localization method using the maximum consensus criterion, leading to a globally (within search radius) optimal solution, and (iii) an assessment of the effects of scene geometry on the distribution of the maximum consensus score.

2. APPROACH

In this part, the localization approach is introduced on exemplary Velodyne VLP-16 point clouds (hereafter referred to as *car sensor* scans). In the evaluation part (section 4), the localization results for a whole trajectory are presented. The used data for the evaluation is introduced in section 3.

For the localization of the vehicle it is assumed that a rough initial *pose* (position and orientation) is available, where the origin of this approximate localization solution is irrelevant. It can be a global navigation satellite system (GNSS) solution, but also be provided by any other kind of initialization, measurement, or prediction from a previous pose. Based on this pose, the car sensor scan can be transformed into the ‘global’ coordinate system (for the purpose of our investigations, we use UTM coordinates), where it should match the globally georeferenced map point cloud. For the purpose of this paper, we use a 3D

pose, but we limit the parameters to be estimated to the (x, y) position in the plane and the heading angle, for a total of three degrees of freedom. Estimation of the height component z is usually less problematic since it can be assumed that the vehicle is moving on the road surface. In addition, it is usually easy to identify and match points on the road surface. As for the angles, roll and pitch are typically very small under normal driving conditions and may also be measured quite reliably.

It is assumed that the true position is within a certain distance from the initial position, and this defines the search range. It is important that this range is chosen large enough to certainly contain the true position, and existing uncertainty information may be used to define it. In theory, it could be extended to cover the entire world, but when considering a meaningful application and acceptable run times, much smaller ranges will usually be used. The task is then to find the position within this range, which maximizes the consensus with regard to all available measurements.

To find this position, the space of possible values (positions) is discretized, and the consensus set is computed for every cell. The consensus count in each cell is defined by the total number of car sensor scan points which match a map point, for that particular position. A match takes place if points are within the distance defined by the edge length of the discretization raster.

The maximum consensus optimization criterion is defined in equation 1. Given the car sensor scan, consisting of the set of points $\{\mathbf{p}_i\}_{i=1}^n$, $\mathbf{p}_i \in \mathbb{R}^3$, and the map point cloud $\{\mathbf{q}_j\}_{j=1}^m$, $\mathbf{q}_j \in \mathbb{R}^3$, the consensus function with respect to the point cloud registration is defined in equation 2, where $\mathbb{I}(c)$ denotes the well-known indicator function which returns +1 if condition c is true, and 0 otherwise. In this investigation, it is sought for a three degrees of freedom rigid transformation $\mathbf{R}_\theta \mathbf{p} + \mathbf{t}$, where the rotation matrix $\mathbf{R}_\theta \in SO(3)$ considers only the heading angle θ and the translation vector $\mathbf{t} \in \mathbb{R}^3$ only the (x, y) translation in the plane, so that $\mathbf{t} = (x, y, 0)^\top = (\mathbf{t}_{xy}, 0)^\top$, with $\mathbf{t}_{xy} \in \mathbb{R}^2$. More precisely, θ and \mathbf{t}_{xy} are discretized, so that each parameter combination corresponds to a cell of the (search space) discretization raster. The inlier threshold ϵ equals the cell size. The outcome of the optimization are the parameters θ^* and \mathbf{t}_{xy}^* , which maximize equation 1.

$$\theta^*, \mathbf{t}_{xy}^* = \arg \max_{\theta \in [-\pi, \pi], \mathbf{t}_{xy} \in \mathbb{R}^2} \Psi(\mathbf{R}_\theta, \mathbf{t}) \quad (1)$$

$$\Psi(\mathbf{R}_\theta, \mathbf{t}) = \sum_{i=1}^n \sum_{j=1}^m \mathbb{I}(\|\mathbf{R}_\theta \mathbf{p}_i + \mathbf{t} - \mathbf{q}_j\|_2 \leq \epsilon) \quad (2)$$

where $\Psi(\cdot)$ is the consensus function,
 $\mathbb{I}(\cdot)$ is the indicator function,
 θ is the heading angle,
 $\mathbf{R}_\theta \in SO(3)$ is the corresponding rotation matrix,
 $\mathbf{t} \in \mathbb{R}^3$ is the translation vector, $\mathbf{t} = (\mathbf{t}_{xy}, 0)^\top$
 $\mathbf{p}_i \in \mathbb{R}^3$ is a point of the car sensor scan,
 $\mathbf{q}_j \in \mathbb{R}^3$ is a point of the map point cloud, and
 ϵ is the inlier threshold.

In this work, we set the size of the overall search space around the initial position to $2m \times 2m$ (from $-1m$ to $+1m$ for both axes) and the edge length of a grid cell to $2cm$ resulting in 100×100 cells within the translation discretization raster. Since

the ‘hits’ are summed up in the discretized space of positions, this is also called the *accumulator*. Technically, instead of iterating over all possible positions, and counting the matches for each position, the accumulator can also be filled by using stencil operations, where for each car sensor scan point, the local environment from the map is added to the accumulator. To avoid extensive zero operations, which occur since most of the space is not occupied, sparse representations can be used. All map points belonging to a local environment can be efficiently selected using a k-d tree spatial data structure. Note that these implementation details increase the efficiency, but do not alter the results in any way. Since points belonging to the ground plane do not create any additional value for the localization in (x, y) and rather interfere with the goal of obtaining a clear peak within the accumulator, ground points are removed in the car sensor scan as well as in the map point cloud beforehand. This is done by discarding all points which have a z -component of the local normal vector close to ± 1 .

Figure 1 visualizes the maximum consensus accumulator for the point cloud depicted in Figure 2. The image shows that the highest consensus set sizes are around the origin $(0, 0)$, which corresponds to the initial position. This is because in this example, the car sensor scans are transformed into the global UTM coordinate system using the reference trajectory on which the map point cloud is based as well. Therefore, the car sensor scan and the map fit perfectly, resulting in the high peak at $(0, 0)$. Apart from the peak, two ‘rays’ can be identified within the maximum consensus grid. Those arise from the structure of the scene. As can be seen in Figure 2, the scene consists of groups of parallel facades, each group defining a plane. Since the normal vectors of the planes point in various directions, the overall set of planes constrains the pose very well, resulting in the clear peak at $(0, 0)$. In addition, each single plane constrains the possible poses only up to a shift parallel to the plane, resulting in rays in the accumulator, which intersect in the origin.




Figure 1. Maximum consensus accumulator for the point cloud depicted in Figure 2.

To investigate the influence of the heading θ , the car sensor scans are rotated around the up-axis in discrete angular steps and the 2D ‘position’ accumulator is calculated for each fixed angle as described above. The 3D accumulator ‘cube’, describing the consensus counts for the full pose, is then obtained by simply stacking the set of 2D accumulators. Figure 3 shows an example for set of position accumulators for different heading angles, which, when stacked, yield the 3D pose accumulator.

One can observe that the heading of -1 degrees leads to the highest and sharpest peak. With increasing deviation from this




Figure 2. Top (upper) and perspective (lower) view of an example point cloud for a crossing scene.




Figure 3. Maximum consensus accumulators for headings from -2 to $+2$ degrees around the initial heading.

heading, the peak gets smaller and more blurred.

In Figure 1, the accumulator is shown for a point cloud obtained at an intersection. As described, due to the structure of the scene, the peak is sharp and unique. In contrast, Figure 4 shows a point cloud in a street without any side roads or crossings. Here, the parallel facades on both sides are the dominant structure in the scene and therefore, the accumulator contains a dominating ray (cf. Figure 5).




Figure 4. Top (upper) and perspective (lower) view of an example point cloud for a straight street without side roads.




Figure 5. Maximum consensus accumulator for the point cloud depicted in Figure 4.

At first sight, the accumulator is characterized by a single ray. However, by searching for its highest peak, it also appears at (0,0). Nevertheless, due to multiple grid cells (along the dominating ray) showing high consensus values, the certainty of the localization result is low compared to the crossing scene depicted in Figure 2. In the evaluation part (section 4), a deeper view on assessing the localization uncertainties is taken.

3. DATA

In this part, the measurement scenario is introduced. Furthermore, the car sensor scans as well as different map data sets, which originate from this measurement, are presented in more detail. Afterwards, the evaluation of the localization approach is conducted based on these different data sets.




Figure 6. Trajectory of measurement experiment (background map: OpenStreetMap (OpenStreetMap contributors, 2021)).

The measurement data was acquired in an inner city area characterized by a dense building structure. The trajectory is shown in Figure 6. The circular path was driven five times.

The car sensor scans were acquired by a Velodyne VLP-16 (Velodyne LiDAR, Inc., 2016). It provides 16 scan layers in a 360° horizontal field of view. The vertical field of view is from -15° to $+15^\circ$, with a uniform vertical spacing of 2° between the layers. The map data is based on point clouds gathered by a Riegl VMX-250 Mobile Mapping System (MMS) (RIEGL Laser Measurement Systems GmbH, 2012). Two Riegl VQ-250 laser scanners are measuring with a rate of up to 300,000 points per second each and a ranging accuracy of ten millimeters. The reference trajectory is acquired using an Applanix POS LV GNSS/IMU system. In a post-processing, by applying correction data from the Satellite Positioning Service (SAPOS), a highly accurate trajectory is obtained.

To obtain more accurate map data, the acquired point clouds of the five runs within the measurement drive are aligned using the adjustment strategy described in (Brenner, 2016). The standard deviation of the resulting map data is below two centimeters. The shifts of the reference trajectory emerging from this adjustment are applied as corrections to the original trajectory. This leads to a highly accurate fit between the adjusted map point cloud and the reference trajectory.

In the evaluation of the localization approach different map data sets are investigated. Those maps are characterized by different point densities referred to as Level Of Detail (LOD). The original adjusted map point cloud corresponds to LOD 0, whereas in this work, due to data handling reasons, investigations are done beginning with LOD 2. The point density between two consecutive LODs is uniformly reduced by a factor of 4. Consequently, the point density of LOD 2 compared to LOD 0 is already reduced by a factor of 16. Figure 7 and 8 show a segment of the map point cloud exemplary for LOD 2 and LOD 4.




Figure 7. Map point cloud in LOD 2.



Figure 8. Map point cloud in LOD 4.

4. EVALUATION

For the evaluation, all accumulators for one measurement drive of the trajectory depicted in Figure 6 are calculated. It comprises 2975 epochs. To assess the localization results obtained by our search algorithm we first analyse the (x, y) positions of the highest consensus sets in the accumulator. According to our premise that the car sensor scans are transformed into global UTM coordinates based on the reference trajectory, we expect the highest peak for each epoch to be close to $(0, 0)$. This expectation is reinforced by the fact that the car sensor scans and the map point cloud were measured simultaneously technically allowing only deviations due to calibration errors. After analysing the (x, y) positions of the highest consensus peaks, we evaluate by which heading angle the highest consensus sets are obtained. Analogous to the position, we expect the deviation from the initial heading to be zero as well. Apart from considering the highest consensus set in the accumulator, we want to use the gained entire knowledge about the search space to evaluate the quality of the localization result. Therefore, we introduce an approach to interpret the distribution of high consensus scores within the accumulator as a measure for certainty.

Before investigating the localization results, it is to mention that all results are presented with respect to the car body frame. The x -axis points across-track to the right of the car, the y -axis points in driving direction and the z -axis points up. In addition to that, the evaluations have shown that there are no significant differences between the four map point clouds with varying LODs. Therefore, only the results using the LOD 4 map point cloud are presented.

In Figure 9, the distribution of the highest peaks over the whole trajectory is depicted in the (x, y) plane. The regular circular pattern which especially appears close to $(0, 0)$ occurs because of the discrete accumulator cells. The mean value of all highest




Figure 9. Qualitative representation of the positions of the highest consensus sets in the 2D (x, y) plane.




Figure 10. Distribution of the highest consensus sets along the x - and y -axes.

peaks is at $(-0.004m, -0.004m)$ and the standard deviations are $\sigma_x = 0.028m, \sigma_y = 0.041m$. Those values indicate that the expected localization position is accurately obtained for the majority of epochs. This perception is encouraged by Figure 10 which shows the quantitative distributions of the highest consensus sets along the x - and y -axes. The highest consensus sets are in most epochs at $(0, 0)$ or at least close to it.

The analysis of the heading angle has shown that it can be determined robustly as well. Figure 11 shows that in most epochs the initial heading angle leads to the highest consensus sets and that in general, the deviation is mostly within $\pm 1^\circ$.

Now that the estimation of the (x, y) -position as well as the heading angle using the highest consensus set in the accumulator provides promising results, we also want to assess the localization quality. To do this, we consider all consensus sets, which reach at least a certain percentage of the score of the largest consensus set. For the evaluation in this work, we set this threshold to 80 percent. We characterize the distribution of the consensus set by computing the mean and covariance of all sets exceeding this threshold. This is graphically represented by an error ellipse centered at the mean. Figure 12 and 13 show the calculated error ellipses, superimposed on the accumulators for the point clouds introduced in Figure 2 and 4.




Figure 11. Distribution of heading angles leading to the highest consensus sets.




Figure 12. For the point cloud depicted in fig. 2: Accumulator and error ellipse based on the weighted sum of all consensus sets reaching at least 80 percent of the highest consensus score.




Figure 13. For the point cloud depicted in fig. 4: Accumulator and error ellipse based on the weighted sum of all consensus sets reaching at least 80 percent of the highest consensus score.

Already by means of those two examples it can clearly be seen that the different scenes inevitably influence the distribution of high consensus sets within the accumulator and therefore the certainty about the localization result. Whereas in the crossing scenario the covariance ellipse is small indicating a comparatively certain localization result, it is small in x -, but large in y -direction for the straight street scenario indicating a high uncertainty in driving direction.

The distribution of the mean values over the whole trajectory is depicted in Figure 14. Compared to the positions of the

highest consensus sets in the (x, y) plane depicted in Figure 9 it is visible that in driving direction (along the y -axis) some mean values significantly differ from the position of the highest consensus set.




Figure 14. For the 2D (x, y) plane: Qualitative representation of the mean values of the consensus sets exceeding the 80 percent threshold.

Analogous to Figure 10, the distribution of the mean values along the x - and y -axes is visualised. The histograms show that for the majority of epochs the mean values also match the initial position at $(0, 0)$.




Figure 15. Distribution of the mean values of the consensus sets exceeding the 80 percent threshold along the x - and y -axes.

In Figure 16 the distribution of the variances of the consensus sets exceeding the threshold is visualized. As the examples in Figure 12 and 13 already indicated, the highest variances occur along the y -axis since in straight streets without any side roads (as they appear for the most time within the trajectory) no single strong peak exists in the accumulator, but rather a line of strong peaks. The facades parallel to the streets are certainly optimal means for robust localization in x -direction, but they also dominate the accumulator due to the high amount of points representing those structures. Varying point densities along the parallel facades may induce false peaks, which outvote the correct solution when orthogonal facade structures are scarce.




Figure 16. Distribution of the variances of the consensus sets exceeding the 80 percent threshold: The highest uncertainties appear in driving direction (along the y -axis).

In Figure 17 and 18 the trajectory and a section of it are visualized along with the error ellipses (enlarged representation) of every 100th (Figure 17) and every 20th (Figure 18) epoch. The different localization qualities for the different moments along the trajectory become visible. In Figure 18, the 2D trajectory position corrected by the shift according to the position of the highest consensus set is visualized as well.

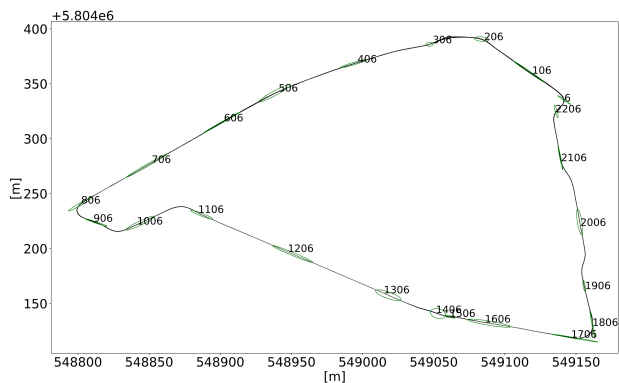


Figure 17. Trajectory of the measurement along with the covariance ellipses of every 100th epoch (ellipses enlarged by a factor of 100).

To conclude the evaluation of the localization quality, the variances along the x - and y -axes over all epochs are visualized in Figure 19. In particular the epochs around number 1800 to 1850 stand out due to their high variances in y -direction. Apart from these few peaks, variances are generally much smaller, however, the overall characteristic that the variances in y are larger than in x holds for most of the epochs.

5. CONCLUSIONS

In this work, we have shown that using maximum consensus based on LiDAR provides great potential for high integrity loc-

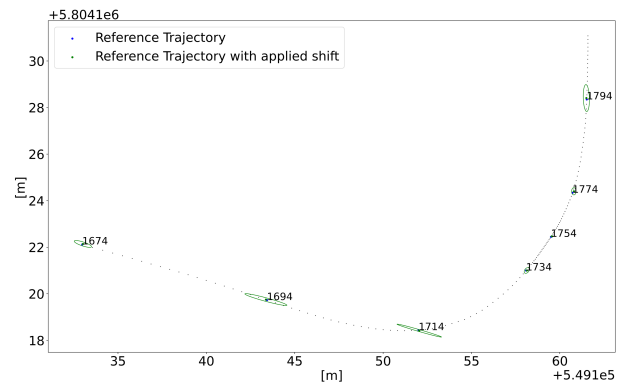


Figure 18. Section of the trajectory along with the positions corrected by the shift according to the highest consensus set and the covariance ellipses of every 20th epoch (ellipses enlarged by a factor of 10).




Figure 19. Variances along the x - (in blue) and y -axes (in green) over all epochs of the measurement.

alization. We have introduced a localization strategy using the maximum consensus criterion based on the registration of two point clouds to obtain a globally optimal 3D pose. In addition to that, tackling the problem by an exhaustive search strategy has given us the opportunity to examine the structure of the search space. This revealed the influence of the scene geometry on the accumulator and offered the possibility to assess the localization results based on the distribution of high maximum consensus scores.

The results also revealed weaknesses of this approach, especially with regard to the along-track localization in street canyons. Tackling those shortcomings is the next step. The idea is to enhance the accumulator by a score which values the distribution of normal vectors within the scene to strengthen possible positions which not only achieve high consensus scores by parallel facades but also perpendicular facing objects. The goal of this approach is to reduce the uncertainty especially in driving direction. Apart from that, the approach will be extended to six degrees of freedom estimating the entire vehicle pose. To limit the computational effort, the application of Branch-and-Bound approaches for this localization strategy will be investigated. A bound on the possible maximum consensus score may also be estimated for each given (map) scene beforehand, based on the scene geometry and a sensor model.

ACKNOWLEDGEMENTS

This project is supported by the German Research Foundation (DFG), as part of the Research Training Group i.c.sens, GRK 2159, 'Integrity and Collaboration in Dynamic Sensor Networks'.

REFERENCES

- Aiger, D., Mitra, N. J., Cohen-Or, D., 2008. 4-Points Congruent Sets for Robust Pairwise Surface Registration. *ACM Trans. Graph.*, 27(3), 1–10. <https://doi.org/10.1145/1360612.1360684>.
- Besl, P. J., McKay, N. D., 1992. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239–256.
- Brenner, C., 2016. Scalable estimation of precision maps in a mapreduce framework. *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, SIGSPACIAL '16, Association for Computing Machinery, New York, NY, USA.
- Breuel, T. M., 1992. Fast recognition using adaptive subdivisions of transformation space. *Proceedings 1992 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 445–451.
- Chen, C.-S., Hung, Y.-P., Cheng, J.-B., 1999. RANSAC-Based DARCES: A New Approach to Fast Automatic Registration of Partially Overlapping Range Images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(11), 1229–1234. <https://doi.org/10.1109/34.809117>.
- Chen, S. Y., 2012. Kalman Filter for Robot Vision: A Survey. *IEEE Transactions on Industrial Electronics*, 59(11), 4409–4420.
- Chetverikov, D., Svirko, D., Stepanov, D., Krsek, P., 2002. The trimmed iterative closest point algorithm. *Object recognition supported by user interaction for service robots*, 3, 545–548 vol.3.
- Chin, T., Cai, Z., Neumann, F., 2018. Robust Fitting in Computer Vision: Easy or Hard? *CoRR*, abs/1802.06464. <http://arxiv.org/abs/1802.06464>.
- Chin, T.-J., Parra Bustos, Á., Brown, M. S., Suter, D., 2014. Fast rotation search for real-time interactive point cloud registration. *Proceedings of the 18th Meeting of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '14, Association for Computing Machinery, New York, NY, USA, 55–62.
- Chin, T.-J., Suter, D., 2017. The Maximum Consensus Problem: Recent Algorithmic Advances. *Synthesis Lectures on Computer Vision*, 7, 1–194.
- Chin, T., Purkait, P., Eriksson, A., Suter, D., 2017. Efficient Globally Optimal Consensus Maximisation with Tree Search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4), 758–772.
- Enqvist, O., Ask, E., Kahl, F., Åström, K., 2012. Robust fitting for multiple view geometry. A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, C. Schmid (eds), *Lecture Notes in Computer Science (Computer Vision - ECCV 2012, Proceedings of the 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Part I)*, 7572, Springer, 738–751.
- Fischler, M. A., Bolles, R. C., 1981. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Commun. ACM*, 24(6), 381–395. <https://doi.org/10.1145/358669.358692>.
- Kalman, R. E., 1960. A New Approach to Linear Filtering and Prediction Problems. *Transactions of the ASME—Journal of Basic Engineering*, 82(Series D), 35–45.
- Le, H., Chin, T., Eriksson, A. P., Suter, D., 2017. Deterministic Approximate Methods for Maximum Consensus Robust Fitting. *CoRR*, abs/1710.10003. <http://arxiv.org/abs/1710.10003>.
- Li, H., 2009. Consensus set maximization with guaranteed global optimality for robust geometry estimation. *2009 IEEE 12th International Conference on Computer Vision*, 1074–1080.
- Li, H., Hartley, R., 2007. The 3d-3d registration problem revisited. *2007 IEEE 11th International Conference on Computer Vision*, 1–8.
- Olsson, C., Enqvist, O., Kahl, F., 2008. A polynomial-time bound for matching and registration with outliers. *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8.
- OpenStreetMap contributors, 2021. <https://www.openstreetmap.org>.
- Parra Bustos, Á., Chin, T., Eriksson, A., Li, H., Suter, D., 2016. Fast Rotation Search with Stereographic Projections for 3D Registration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(11), 2227–2240.
- RIEGL Laser Measurement Systems GmbH, 2012. RIEGL VMX-250. Technical report.
- Velodyne LiDAR, Inc., 2016. VLP-16 Velodyne LiDAR Puck.
- Yang, J., Li, H., Jia, Y., 2013. Go-icp: Solving 3d registration efficiently and globally optimally. *2013 IEEE International Conference on Computer Vision*, 1457–1464.
- Zheng, Y., Sugimoto, S., Okutomi, M., 2011. Deterministically maximizing feasible subsystem for robust model fitting with unit norm constraint. *CVPR 2011*, 1825–1832.