

A COMPARISON OF METHODS FOR LINE EXTRACTION FROM RANGE DATA

Daniel Sack * Wolfram Burgard *

** University of Freiburg, Dept. of Computer Science, Germany*

Abstract: The representation of the environment of a mobile robot by line models is a popular alternative to occupancy grid maps. Line maps require significantly less memory than occupancy grids and therefore scale better with the size of the environment. They furthermore are more accurate since they do not suffer from discretization problems. In the past a variety of techniques for learning line maps from range data have been developed. These techniques differ in various aspects such as the way lines are extracted from range scans, how the lines are updated upon sensory input. There furthermore are techniques that are able to operate online, whereas others postprocess the data. In this paper we compare three different techniques for learning line models with respect to various parameters such as efficiency and quality of the resulting maps. Experimental results illustrate the advantages and the disadvantages of the different techniques.

Keywords: Line models, mobile robot mapping

1. INTRODUCTION

Geometric maps play an important role in mobile robotics since they support various tasks such as path planning and accurate localization. One of the most popular way to represent the environment of a robot are occupancy grids. Their advantage is that they can easily be updated upon sensory input. Their disadvantages, however, lie in the huge memory requirements and the the limited accuracy due to the discretization. To overcome these limitations, many authors studied the generation of more compact representations such as line models. Line maps are based on the assumption that most indoor environments consist of planar structures such as walls, doors, or cupboards. They are more compact since such objects can effectively be represented by a small number of line segments whereas thousands of grid cells might be required by an occupancy grid. They are also more accurate since they provide floating point resolution and do not suffer from discretization problems.

Accordingly techniques for extracting line models from range data have been studied intensively in the

past. Crowley (1989) extracts line segments from range measurements and combines these segments using a Kalman filter. Pfister et al. (2003) extend this approach and also consider the accuracy of the measurements when updating the line model. Arras and Siegwart (1997) use a clustering approach to learn line models from laser data. Their approach also considers the uncertainty in the measurements when clustering points into linear segments. The approach developed by Gonzales et al. (1994) computes point clusters from each range scan based on the distance between consecutive points. They apply linear regression to fit lines to these clusters and iteratively combine lines to a global map. Leonard et al. (2001) use a Hough transform to extract linear features from sequences of consecutive sonar measurements. These features are then maintained using a Kalman filter. Several approaches apply the well-known iterative end-point fit or split-and-merge algorithm developed by Duda and Hart (1973) for fitting lines to scans. Schröter et al. (2002) adopt the approach of Gutmann et al. (2001) to cluster scans using the split-and-merge algorithm. They combine nearby segments using a weighted vari-

ant of linear regression. Also Newmann et al. (2002) use this approach in combination with the Ransac algorithm proposed by Fischler and Bolles (1981) to extract linear models from laser data. Baltzakis and Trahanias (2002) propose an approach for simultaneous localization and mapping in which the features are extracted using the split-and-merge algorithm. Recently Liu et al. (2001) applied the EM-algorithm to extract planar structures from 3d data. An online variant of this approach has been presented by Martin and Thrun (2002).

Whereas all these approaches seek to find a minimum number of line segments that best approximate a given set of range data, the underlying techniques are quite different. Some of the techniques use the split-and-merge algorithm, some rely on the Hough transform, and others use the EM-algorithm for clustering data points into lines. Furthermore, some of the approaches have been designed to operate online, i.e., while the robot is exploring its environment. Others, in contrast, are offline techniques that are run after the robot has acquired all relevant information. In this paper we study three different approaches for line extraction from laser range data. The first approach is an incremental online technique similar to the approaches of Crowley (1989), Schröter et al. (2002), and Pfister et al. (2003). The second method is an offline technique, which relies on ideas similar to the online techniques but operates on the whole data set to compute linear approximations. The third approach is a variant of the EM technique proposed by Liu et al. (2001). Throughout this paper we analyze and compare these approaches with respect to various features such as computational requirements, accuracy, and also robustness. This paper is organized as follows: The next section describes the techniques compared in this paper. Section 3 then summarizes the comprehensive experiments carried out with various data sets generated with real robots and in simulation.

2. APPROACHES TO LEARN LINE MODELS

2.1 Incremental Learning of Line Models

The approach to incrementally learning line models analyzed here is closely related to the approach developed by Schröter et al. (2002) and borrows several ideas from the technique proposed by Pfister et al. (2003). The key idea is to extract line segments from laser range scans and to integrate these new lines with the line segments stored in a global map.

There are two popular approach to extract lines from individual range scans. The Hough transform employed by Pfister et al. (2003) is typically used to cluster collinear points, which then are approximated by lines. The split-and-merge algorithm, on the other hand, recursively subdivides the scan into sets of neighboring beams that can accurately be approxi-

mated by lines. The split-and-merge algorithm has several desirable properties compared to approaches using the Hough transform. It also considers the order in which the measurements were obtained, it typically runs much faster than algorithms that first have to compute the Hough transform, and it exploits the local structure whereas the Hough transform only considers the global structure. Also, the results generated by the split-and-merge algorithm are locally more consistent than those obtained with the Hough transform.

Therefore our approach for incrementally learning line models uses the split-and-merge algorithm to extract line segments from range scans and afterwards combines these line segments with a given global map. To fit a line to a set of points we use the approach proposed by Lu and Milios (1994). In this approach the line that best approximates a set of points is computed according to

$$\tan 2\phi = \frac{-2 \sum_i (\bar{x} - x_i)(\bar{y} - y_i)}{\sum_i ((\bar{y} - y_i)^2 - (\bar{x} - x_i)^2)} \quad (1)$$

$$r = \bar{x} \cos \phi + \bar{y} \sin \phi \quad (2)$$

where $\bar{x} = \frac{1}{n} \sum_i x_i$, $\bar{y} = \frac{1}{n} \sum_i y_i$, r is the normal distance of the line from the origin, and ϕ is angle of the normal.

To integrate a line segment s_l extracted from a single scan into a map m consisting of line segments we search for the line segment s^* in m that is closest to s_l . The distance between s_l and s^* is given by the sum of the distances of the endpoints of s_l from the line going through s^* . We combine s_l and s^* only when the distances of both endpoints do not exceed a certain threshold, which is 10cm in our current system, and if s_l partly overlaps with s^* . To compute a new line from the two segments s_l and s^* we combine the data points associated to the two line segments and compute a new line from the resulting point set using Equations 1 and 2.

Please note that using this approach the complexity of combining a new line segment with a line stored in m is $O(nk)$ in the worst case where n is the number of scans taken so far and k is the number of range measurements per scan. To achieve a constant complexity per measurement one can sub-sample the data points. Alternatively one can weigh the segments according to their length (such as applied by Schröter et al. (2002)), according to the number of data points used to compute the individual line segments, or according to the variances of the line parameters (see Crowley (1989) and Pfister et al. (2003)). In our current system we used a sub-sampling approach in which we limited the number of range measurements to 20,000 which turned out to be efficient enough for an online integration of new lines. Note that this number was rarely reached in any of our experiments. Furthermore we never observed any evidence that the sub-sampling lead to significantly less accurate results.

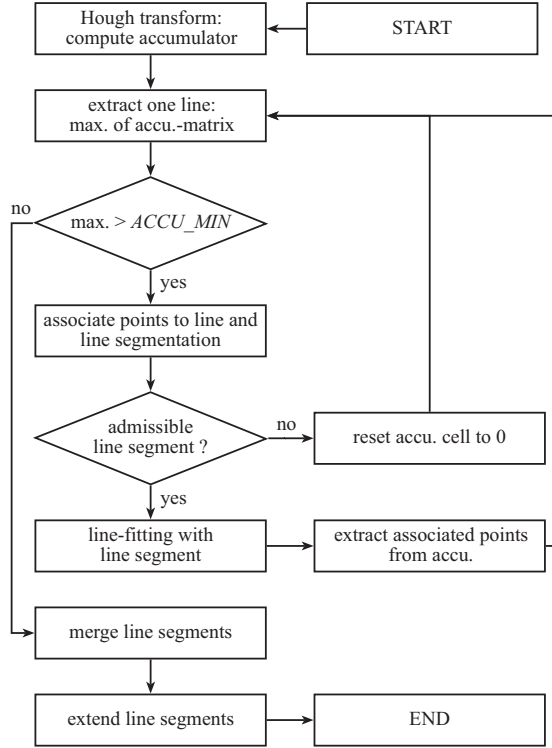


Fig. 1. Flow-Chart for the offline approach.

2.2 An Offline Approach for Learning Line Models

Online techniques such as the one described above have the advantage that they can be applied while the robot is moving through the environment. The disadvantage of such approaches is that it is harder to detect and filter out dynamic aspects such as beams reflected by people walking by and outliers. If one does not use features for the detection of corrupted readings, one in principle has to consider all past data during the map learning process. This, however, introduces a complexity logarithmic in the size of the past data for updating the map based on a single scan.

The approach implemented for a global approximation of range data by line segments is summarized in Figure 1. In a preprocessing step we compute the accumulator of the Hough transform for all scan points. We then iterate the following procedure on the accumulator until the value of the maximum falls below a certain threshold ($ACCU_MIN$): We take the maximum of the accumulator and compute the points associated or close to the line that corresponds to this maximum. In the next step we compute the line segments for the extracted line and take the segment with the maximum number of associated points. If the number of points for this segment is too small, we simply set the corresponding accumulator cell to zero. Otherwise we fit a line to the points associated to that segment and delete the points from the accumulator of the Hough transform. As a result the number of points in the accumulator steadily decreases until $ACCU_MIN$ is no longer exceeded by the maximum value. Note that this approach iteratively finds the greatest local maxima in

the accumulator. Thereby it adapts the accumulator according to the linear structures extracted so far. When the loop has terminated we merge all line segments using the same algorithm as in the incremental approach but without sub-sampling the points. In a final step we extend the length of line segments to connect neighboring line segments.

To deal with dynamic objects we use a simplified version of Hähnel et al. (2003). We compute an occupancy grid map of the environment and identify all beams that end in cells not covered by obstacles. These beams are regarded as reflected by dynamic objects and filtered out before the line extraction process described above. Note that this approach is able to remove readings reflected by persons walking by. However, it is unable to correctly model objects that are moved around as in the approach of Angelou et al. (2002).

2.3 Learning Line Models with EM

One popular approach for learning line models from laser range data is using the EM algorithm (see Liu et al. (2001)). Throughout this paper we use a variant of the EM algorithm also known as fuzzy k-means clustering (see Duda and Hart (1973)). This approach assumes a uniform noise model for all beams obtained from the laser range sensor. Accordingly, the likelihood of a measurement z given a line θ is defined as

$$p(z | \theta) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2} \frac{d^2(z, \theta)}{\sigma^2}} \quad (3)$$

The goal of the EM algorithm is to generate an iterated sequence of models of increased likelihood. To achieve this, one introduces so-called correspondence variables that specify which measurement belongs to which linear component of the model. Since the correct values of these assignment variables are unknown, one estimates a posterior about the value of these correspondence variables. Let θ_j be a component of the model and z_i be a measurement. Then the expectation about c_{ij} , i.e., that measurement i belongs to line j is computed in the so-called E-Step as

$$E[c_{ij} | \theta, z] = p(c_{ij} | \theta, z) \quad (4)$$

$$= \alpha p(z | c_{ij}, \theta) p(c_{ij} | \theta) \quad (5)$$

$$= \alpha' p(z_i | \theta_j). \quad (6)$$

In the M-step, the algorithm then computes the parameters of the model by taking into account the expectations computed in the E-step:

$$\theta^* = \operatorname{argmin}_{\theta'} \sum_i \sum_j E[c_{ij} | \theta, z] d^2(z_i, \theta'_j) \quad (7)$$

Given a fixed variance over all data points in z we fortunately can determine the most likely model in closed

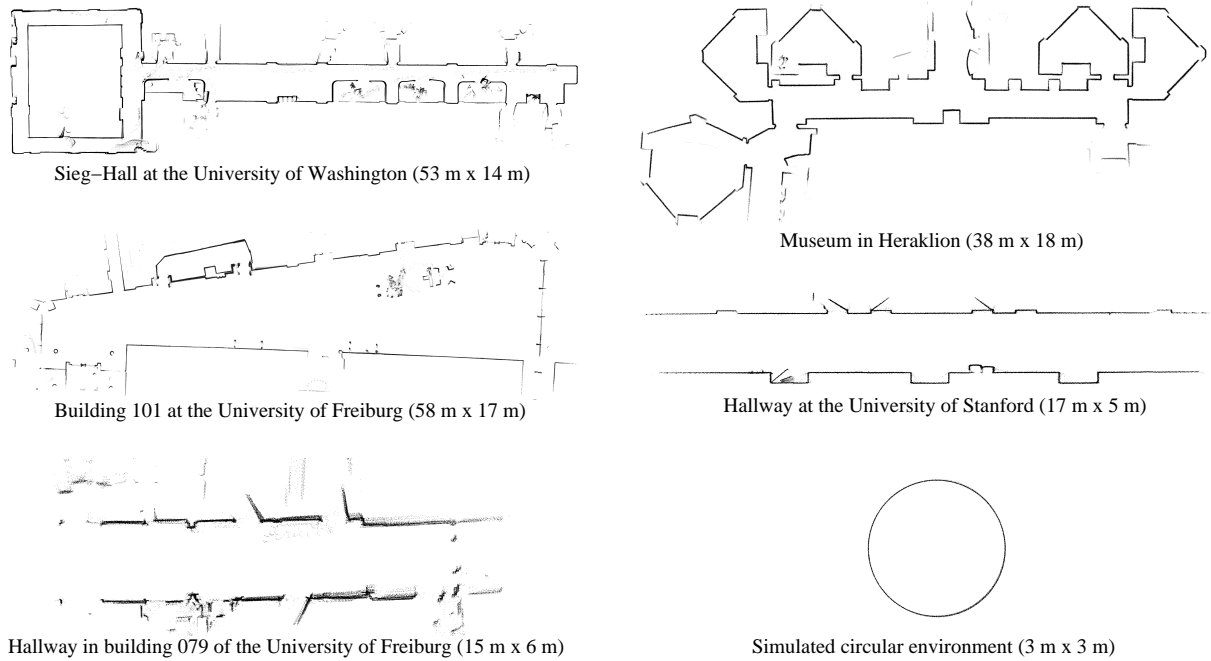


Fig. 2. Point maps of the different environments used in this paper to analyze the algorithms.

form according to the following equations (see Arras and Siegwart (1997)):

$$\tan 2\phi_j = \frac{-2 \sum_i E[c_{ij}|\theta, z](\bar{x} - x_i)(\bar{y} - y_i)}{\sum_i E[c_{ij}|\theta, z]((\bar{y} - y_i)^2 - (\bar{x} - x_i)^2)} \quad (8)$$

$$r_j = \bar{x} \cos \phi_j + \bar{y} \sin \phi_j \quad (9)$$

Here \bar{x} and \bar{y} are computed as

$$\bar{x} = \frac{\sum_i E[c_{ij}|\theta, z]x_i}{\sum_i E[c_{ij}|\theta, z]}, \quad \bar{y} = \frac{\sum_i E[c_{ij}|\theta, z]y_i}{\sum_i E[c_{ij}|\theta, z]} \quad (10)$$

A crucial problem when applying EM is the number of model components. The algorithm described here applies the approach also used by Bennewitz et al. (2002). Whenever the algorithm has converged we remove lines that have a low utility, i.e., lines that can safely be discarded. Additionally, we consider data points with a low support and introduce new lines for such points. The final solution is that model that yields the maximum likelihood over all data points.

3. EXPERIMENTAL RESULTS

The techniques described above have been implemented and evaluated using several laser range data sets gathered with mobile robots. The goal of the experiments described here is to analyze the different approaches with respect to several properties, such as the computational requirements, the number of line segments extracted, the ability to approximate small planar or non-linear objects, and the robustness against to spurious measurements.

The number of points in the data sets were 1.8 million for the Sieg-Hall, 671,976 for the building 101

and 35,399 for the building 079 of the University of Freiburg, 283,458 for the Heraklion museum data, 162,554 for the Stanford corridor, and 10,190 for the circular environment.

Figure 2 shows the sets of points we used to analyze the algorithms. Two of the data sets are from larger unstructured environments such as the museum in Heraklion or entrance Hall of building 101. Additionally, we have three data sets from corridor environments with different levels and kinds of sensor noise. Whereas the data set from the corridor of the Stanford University is quite accurate, the data gathered at University of Washington contain several data points generated by dynamic objects. Finally, the point map of the corridor of building 079 includes noisy measurements due to glass panes. Additionally, we considered simulated data from a circular environment in order to analyze the capability of the techniques to approximate non-planar structures.

Figure 3 left diagram shows the number of line segments generated by the individual approaches for the six different data sets. For the huge data sets (Sieg-Hall and building 101) we aborted the EM algorithm after three days of computing. As can be seen from the plot the number of line segments generated by all three techniques were quite similar. The increased number of segments generated by the incremental approach for the Sieg-Hall data is caused by the huge number of data points generated by people walking in the environment. These data points could be filtered out in the offline and the EM approach.

The runtime of the algorithms is depicted in Figure 3 center diagram. As can be seen from the figure, the online approach always completes the task before the other two methods. Additionally, it is worth noting




Fig. 3. Number of extracted line segments (left) and time to compute the line maps on an Intel Celeron 2.8 GHz computer (center) for the different data sets. Approximation error per point in the resulting line maps (right).




Fig. 4. Fraction of line maps of the museum in Heraklion generated by the different approaches: Incremental approach (left), offline approach (center) and EM-algorithm (right).

that the EM algorithm has an enormous overhead which leads to increased completion times. Note that the scale of y -axis is logarithmic.

Finally, Figure 3 right diagram plots the average distance of a data point from a line segment for the different techniques and data sets. To compute this value we only considered points that were at most 1m away from a line segment. It turns out that the offline techniques yields the best approximations in four of the six cases. For the building 101, the incremental approach was better since resulting model contained 5% more line segments than the model generated by the offline approach. In the circular structure of the simulated data, the EM algorithm produced the best result.

For all three approaches we observed that larger linear structures were discovered quite accurately. The major differences appeared in the approximation of smaller structures in the larger maps. As an example consider the data set of the Heraklion museum. The line map generated by the offline techniques is depicted in Figure 5. Parts of the maps generated by the three different approaches are shown in Figure 4. As can be seen from this figure, the incremental and the offline approaches produces almost the same result (see left two figures). The shorter linear structures were not approximated well using the EM algorithm. In the incremental and offline approach we directly fit segments to data points whereas the EM technique considers lines instead of segments so that often points far away from a linear structure, which is approximated by a line, have an influence on the parameters of that line.

To see the difference between the incremental approach and the offline technique consider Figure 6. As already mentioned above, the point map of the Sieg-Hall data set contains several data points corresponding to dynamic objects. Accordingly, the incremental approach generates a larger number of line segments since it does not include means to filter




Fig. 5. Line map of the museum in Heraklion computed by the offline approach.



Fig. 6. Fraction of the line maps for the Sieg-Hall data set generated by the incremental (left) and offline (right) approach.

out measurements corresponding to dynamic objects. On the other hand, the doorways illustrate that the incremental technique allows a better extraction of small linear structures. This is because the split-and-merge algorithm operates on the individual scans in which these linear structures clearly appear. In the global point map, however, these small linear structures can only hardly be detected because of the noise in the range measurements and slight registration errors.

The final experiment has been carried out using a point map of a circular room generated in simulation. Figure 7 depicts the maps generated by the individual techniques. Whereas all three maps are quite accurate, the EM approach yields the best result. Because the EM algorithm performs a global optimization the lines converge to a regular polygon. The incremental and the offline approach, however, do not perform a global optimization so that the overall result always is suboptimal (see also Figure 3 right diagram).




Fig. 7. Line maps of the circular environment generated by different approaches: incremental approach (left), offline approach (center) and EM-algorithm (right).

Table 1 summarizes the properties of the line mapping techniques discussed in this paper. It turns out that the strength of the incremental technique lies in the efficiency and the capability to extract short linear structures. Both, the offline technique and the incremental technique perform well on large environments, whereas the EM approach fails due to the computational requirements. One advantage of the offline and the EM technique is that they can better deal with dynamic aspects and noise since they can retrospectively identify measurements as dynamic and filter them out whereas online-techniques typically generate short line segments for such data. The strength of the EM approach appears when it comes to the approximation of non-linear structures. Since the EM technique seeks to globally optimize the parameters of a map, the results for circular structures are better than with the other methods. The major disadvantage of the EM approach comes from the consideration of lines instead of line segments. Accordingly, data points far a way from a linear structure can influence their parameters. However, the EM approach can easily be extended to situations in which the uncertainty of the data points is not the same for all points.

4. CONCLUSIONS

In this paper we compared three different approaches for learning line maps from range data. These methods differ in important aspects and show similarities to various approaches found in the literature. In extensive experiments we analyzed the properties of these techniques and identified their advantages and disadvantages. One of the most surprising result is that the EM technique, although it belongs to the global optimization techniques, does not perform better than offline and online techniques that follow a greedy maximization scheme. It furthermore appeared that the online technique, which consider individual scans, generates more accurate results for smaller structures than the offline and the EM techniques.

This opens several interesting directions for future research. One important question is, whether the EM-based approach can be extended to operate on line segments instead of lines. Furthermore, the question of how to consider the individual scans in the offline technique would improve the quality of smaller linear structures especially in large maps. Finally, the filtering of dynamic objects in online approaches appears to be an interesting topic for future research.

Table 1. Summary of the comparison.

	incr.	offline	EM
speed	++	+	--
short segments	+	o	-
large environments	++	+	--
nonlinear objects	o	o	+
dynamic objects	-	+	+
robustness	-	+	o

REFERENCES

- D. Anguelov, R. Biswas, D. Koller, B. Limketkai, S. Sanner, and S. Thrun. Learning hierarchical object maps of non-stationary environments with mobile robots. In *Proc. of the 17th Annual Conf. on Uncertainty in AI (UAI)*, 2002.
- K. Arras and R. Siegwart. Feature extraction and scene interpretation for map-based navigation and map building. In *Proc. SPIE, Mobile Robotics XII*, volume 3210, 1997.
- H. Baltzakis and P. Trahanias. An iterative approach for building feature maps in cyclic environments. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.
- M. Bennewitz, W. Burgard, and S. Thrun. Using EM to learn motion behaviors of persons with mobile robots. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.
- J. Crowley. World modeling and position estimation for a mobile robot using ultrasound ranging. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1989.
- R.O. Duda and P.E. Hart. *Pattern classification and scene analysis*. Wiley, New York, 1973.
- M.A. Fischler and R.C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image and analysis and automated cartograph. *Comm. of the ACM*, 24(6), 1981.
- J. Gonzales, A. Ollero, and A. Reina. Map building for a mobile robot equipped with a 2d laser rangefinder. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1994.
- J.S. Gutmann, T. Weigel, and B. Nebel. A fast, accurate, and robust method for self-localization in polygonal environments using laser-range-finders. *Advanced Robotics*, (8):651–668, 2001.
- D. Hähnel, R. Triebel, W. Burgard, and S. Thrun. Map building with mobile robots in dynamic environments. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.
- J. Leonard, P. Newmann, R.J. Rikoski, J.D. Tardós, and J. Neira. Towards robust data association and feature modeling for concurrent mapping and localization. In *Proc. of the Tenth Int. Symposium on Robotics Research (ISRR)*, 2001.
- Y. Liu, R. Emery, D. Chakrabarti, W. Burgard, and S. Thrun. Using EM to learn 3D models of indoor environments with mobile robots. In *Proc. of the Int. Conf. on Machine Learning (ICML)*, 2001.
- F. Lu and E.E. Milios. Robot pose estimation in unknown environments by matching 2d range scans. In *Proc. of the IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 935–938, 1994.
- C. Martin and S. Thrun. Online acquisition of compact volumetric maps with mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.
- P. Newmann, J. Leonard, J.D. Tardós, and J. Neira. Explore and return: Experimental validation of real-time concurrent mapping and localization. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.
- S.T. Pfister, S.I. Roumeliotis, and J.W. Burdick. Weighted line fitting algorithms for mobile robot map building and efficient data representation. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.
- B. Schröter, M. Beetz, and J.-S. Gutmann. RG mapping: Learning compact and structured 2d line maps of indoor environments. In *Proc. of the Int. Workshop on Robot and Human Interactive Communication (ROMAN)*, 2002.