Maximum-Likelihood Sample-Based Maps for Mobile Robots

Daniel Meyer-Delius   Wolfram Burgard

University of Freiburg, Department of Computer Science, D-79110 Freiburg

{meyerdel|burgard}@informatik.uni-freiburg.de

Abstract— The problem of representing environments of a mobile robot has been studied intensively in the past. The predominant approaches for geometric representations are grid-based or line-based maps. In this paper, we consider sample-based maps which use the data points obtained by range scanners to represent the environment. The main advantage of this representation over the other techniques is that it does not impose any a priori structure on the environment. However, range measurements come in large amounts. We present a novel approach for calculating maximum-likelihood subsets of the data points by sub-sampling laser range data. In particular, our method applies a variant of the fuzzy k-means algorithm to find a map that maximizes the likelihood of the original data. Our approach has been implemented and tested on real data gathered with a mobile robot.

Index Terms— Mobile robots, mapping, likelihood maximization

I. INTRODUCTION

Geometric representations of the environment play an important role in mobile robotics since they support various fundamental tasks such as path planning or localization. One of the most popular approaches are occupancy grids, which provide a discrete probabilistic representation of the environment. Other popular approach are representations based on geometric primitives that are typically found in the environment. In this contexts, lines play a major role, since many man-made buildings are composed of linear structures like walls, for example. Although these techniques have been successfully applied in the past, they have certain disadvantages coming from discretization errors or because of missing features in the environment.

In this paper we consider a non-parametric representation of the environment which uses the data points gathered by the robot itself for the representation. In particular, this point-based representation, which will be denoted as sample-based maps throughout this paper, utilizes the endpoints of the range measurements projected into the global coordinate system to represent the objects in the environment. This representation, which rests on the assumption that the data is its best own model, has several advantages. First it provides a high accuracy at floating point resolution. Additionally, it does not suffer from discretization errors like grid maps. Furthermore it is flexible, as it does not rely on any pre-defined features which can be regarded as a certain a priori information about the environment. From this point of view, sample-based maps inherit the advantages of grid maps and feature-based maps
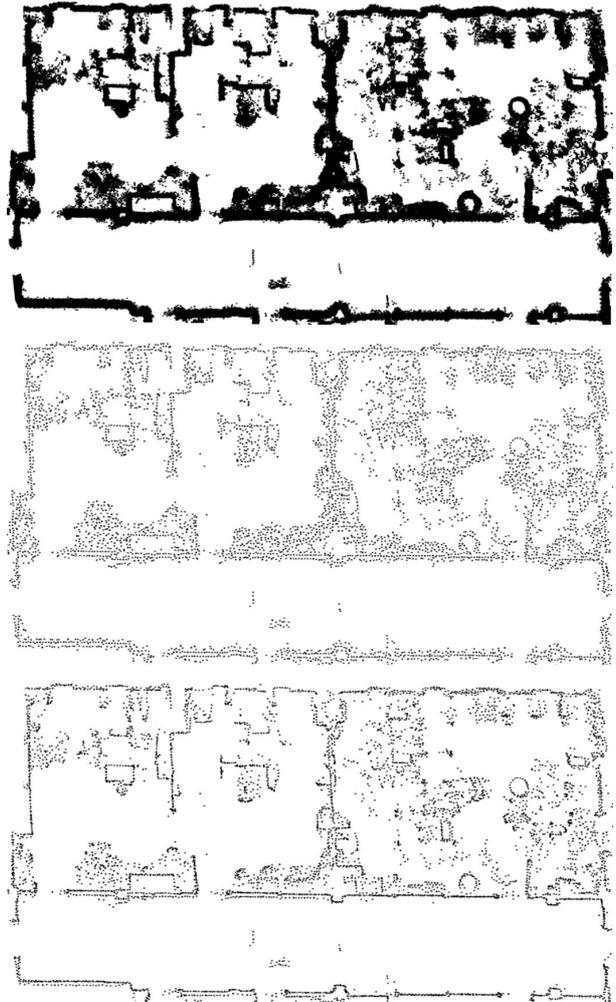
Fig. 1. Typical sample-based maps. The top image depicts a dataset consisting of 706,628 data points obtained in an office building. The map depicted in the middle is obtained by down-sampling the dataset. It consists of less than 1% of the samples in the original dataset. The lower image shows a maximum-likelihood map obtained with our approach with the same number of samples than the map in the middle.

as they are able to represent arbitrary structures and at the same time provide an arbitrary accuracy. Finally, sample-based maps are consistent with the data perceived by the robot and do not require any procedures for transforming them into another internal structure.

One drawback of sample-based maps, however, is that range

measurements come in large amounts which typically introduce enormous memory requirements. At the same time, many of the points provided by a range sensor are redundant because nearby points already provide the relevant information about the existence of an object in the environment. One popular approach to reduce the number of points in a sample-based map is to down-sample the data points. However, this directly introduces the question of how to down-sample, especially because not every sample necessarily contributes to the model in the same way.

As an example, consider the maps depicted in Figure 1. The topmost map shows a complete data set obtained with a Pioneer II robot equipped with a laser range finder in building 79 at the University of Freiburg. The complete dataset consists of 706,528 points. The map in the middle was obtained by down-sampling the whole dataset using grid-based sampling, one of the techniques described in this paper. It contains only 5,749 points, which is less than 1% of the points in the original dataset. Even with this enormous simplification, the structure of the environment is still recognizable. The bottom map contains also 5,749 points and is the result of maximizing the likelihood of the complete dataset using the approach described in this paper. Observe how, as result of the maximization, the structures in the environment are defined more clearly.

In this paper, we consider the problem of finding a subset of data points of a sample-based map that provides an efficient representation of the environment and at the same time maximizes the likelihood of all the data observed by the robot. Our approach applies a variant of the fuzzy k-means algorithm to find a subset that maximizes the likelihood of the whole dataset. This algorithm can be regarded as an instance of the expectation maximization (EM) algorithm and describes an approximative optimization technique that takes an initial map and iteratively improves it in a hill-climbing fashion maximizing the likelihood of the dataset. As EM highly depends on the initialization, we also describe a strategy for choosing the initial point set. As a result, we obtain highly accurate maps which consist only of a small fraction of the original data and at the same time maximize its likelihood. Additionally, the maps generated by our approach improve the capabilities of the robot to localize itself in the environment using a range scanner.

This paper is organized as follows. After discussing related work in the following section, we will describe how the fuzzy k-means algorithm can be applied to find the subset of samples that maximizes the likelihood of the original data. Then, in Section IV, we will present a technique for initializing the fuzzy k-means algorithm. Finally, we will present results obtained with our approach in Section V.

II. RELATED WORK

In the robotic literature many different strategies have been proposed to learn geometric representations from range data. One of the most popular approaches, occupancy grid maps, have been introduced by Moravec and Elfes [12]. Whereas this approach has a nice probabilistic formulation and additionally allows a robot to maintain a map over time, occupancy grid maps require a huge amount of memory and suffer from discretization errors.

Besides occupancy grid maps, line-based representations have been very popular. One of the first such approaches has been described by Crowley [4] who uses a Kalman filter to fit lines on range measurements obtained from sonar sensors. In the method proposed by Gonzalez *et al.* [7], point clusters are computed from each range scan based on the distance between consecutive points. Linear regression is then applied to fit lines to these clusters. Arras and Siegwart [2] use a hierarchical clustering approach to extract lines from the points obtained from laser data. The strategy proposed by Leonard *et al.* [10] uses a Hough transform to extract linear features from a sequence of consecutive sonar measurements. The approach presented by Schröter *et al.* [16] clusters scans using the split-and-merge algorithm and combines nearby segments using a weighted variant of linear regression. Other geometrical primitives have also been used for describing the environment. González-Baños and Latombe [8] use polylines as primitives for the representation. A polyline is a sequence of line segments that are connected at their endpoints. Veeck and Burgard [20] describe an approach for learning polyline maps that operates on an arbitrary set of points. Whereas all these techniques provide highly accurate maps, they rest on the assumption, that the environment actually consists of linear structures.

In contrast to this, our work described in this paper does not extract geometrical structures from the data points, but uses the points themselves as primitives for the representation. In the computer graphics field, Levoy and Whitted [11] proposed points as a universal meta-primitive for geometric modeling and rendering applications for 3D geometry. Pauly *et al.* [13] explore the usage of points as primitives for modeling three-dimensional objects, and present several techniques for modifying and reducing the size of the original set of points. However, these techniques produce set of points that are not necessarily a subset of the original dataset, creating models that can be inconsistent with the data, especially when the generated points lie in free-space. Alexa *et al.* [1] work on a set of samples that is a true subset of the original dataset, but their method relies on the idea that the given set of points defines a surface, imposing in this way an a priori structure. In contrast to this, our work does not make any a priori assumptions on the dataset and uses a subset of the data as its best model. This subset is calculated using a probabilistic formulation of the overall problem and the application of the EM to find a maximum likelihood subset. In the past, Thrun [18] has applied the EM algorithm to find grid maps that maximize the likelihood of the data. In this paper we borrow this idea and calculate subsets of sample points that maximize the likelihood of the original data.

The idea of using samples to represent the environment is not new. Biber and Duckett [3], for example, porpose a dynamic map that adapts continously over time using a sample-based representation of the environment. Their work is on dealing with the dynamics of the environment, while ours is on finding an efficient representation for it.

III. SAMPLE-BASED MAPS

A sample-based map is a non-parametric representation of the environment which uses points to represent the objects in the environment. These points correspond to the endpoints of the range measurements gathered by the robot projected into the global coordinate system. In order to project the measurements, the pose of the robot must be known for all range measurements, or an accurate pose estimation must be given. Projecting all range measurements, a set $\mathcal{D} = \{x_1, \ldots, x_N\}$ of points is obtained. Formally, a sample-based map $\mathcal{S} = \{x_1, \ldots, x_K\}$ is a subset $\mathcal{S} \subseteq \mathcal{D}$ of the dataset $\mathcal{D}$.

A. Probabilistic Interpretation

A sample-based map $\mathcal{S}$ can be interpreted as a probabilistic mixture model Θ with K components, where each point $x_i \in \mathcal{S}$ corresponds to a component ω_i in the model. The following assumptions are made about the probabilistic structure of the model:

1) The number K of components ω_j, that is, the size of the model, is known.
2) Each component ω_j has the same prior probability $P(\omega_j) = 1/K$, for $j = 1, \ldots, K$.
3) Each sample $x_i \in \mathcal{D}$ is independently and normally distributed according to $p(x_i|\omega_j, \theta_j) \sim \mathcal{N}(\mu_j, \Sigma_j)$, where θ_j represents the mean μ_j and covariance matrix Σ_j of the normal distribution associated to component ω_j.

Samples in $\mathcal{D}$ are assumed to be generated by first selecting a component ω_j with probability $P(\omega_j)$ and then selecting a sample x according to $p(x|\omega_j, \theta_j)$. The probability density function for a sample x is then given by the following mixture density

$$p(x|\Theta) \;=\; \sum_{j=1}^{K} p(x|\omega_j, \theta_j) P(\omega_j)\,, \qquad (1)$$

where $\Theta = \{\theta_1, \ldots, \theta_K\}$. We further assume that the covariance matrices Σ_j of the probabilities $p(x|\omega_j, \theta_j)$ are isotropic, known, and identical for each component ω_j. Based on these assumptions, our unknown parameter vector $\Theta = \{\mu_1, \ldots, \mu_K\}$ consists of only the means μ_i of the normal distributions from which the samples are selected.

B. Likelihood Maximization

If $\mathcal{D}$ contains N samples, the likelihood of the data $\mathcal{D}$ given the parameters Θ is computed as

$$p(\mathcal{D}|\Theta) \;=\; \prod_{i=1}^{N} p(x_i|\Theta)\,. \qquad (2)$$

We are interested in finding the subset $\mathcal{S} \subseteq \mathcal{D}$ of points that maximizes the likelihood of the data $\mathcal{D}$. The likelihood is a measure of the goodness of the model. It indicates, for a fixed dataset, that a model with a higher likelihood value is more likely to be the true model than a model with a lower one.

Since evaluating every possible subset $\mathcal{S}$ is computationally infeasible, our approach to find a maximum-likelihood map is to used an approximative iterative optimization technique. We consider a sample-based map $\mathcal{S}$ as a probabilistic model Θ and use a variant of the the fuzzy k-means algorithm for finding the value of Θ that maximize the likelihood of $\mathcal{D}$. The fuzzy k-means algorithm is basically a gradient ascent or hill-climbing procedure which seeks a minimum of the following heuristic cost function [5]

$$J \;=\; \sum_{j=1}^{K}\sum_{i=1}^{N} P(\omega_j|x_i, \mu_j) \|x_i - \mu_j\|^2\,, \qquad (3)$$

where $P(\omega_j|x_i, \mu_j)$ is an estimate for the posterior probability for ω_j defined as

$$
\begin{aligned}
P(\omega_j|x_i, \mu_j) &= \frac{p(x_i|\omega_j, \mu_j) P(\omega_j)}{p(x_i|\Theta)} \\
&= \frac{p(x_i|\omega_j, \mu_j)}{\sum_{l=1}^{K} p(x_i|\omega_l, \mu_l)}\,. \qquad (4)
\end{aligned}
$$

The algorithm uses an initial estimate $\Theta^{[0]} = \{\mu_1^{[0]}, \ldots, \mu_K^{[0]}\}$ to compute the posteriors $P(\omega_j|x_i, \mu_j^{[0]})$ using (4) and then updates the estimates $\Theta^{[0]}$ according to

$$\mu_j^{[1]} \;=\; \frac{\sum_{i=1}^{N} P(\omega_j|x_i, \mu_j^{[0]}) x_i}{\sum_{i=1}^{N} P(\omega_j|x_i, \mu_j^{[0]})}\,. \qquad (5)$$

The improved estimates are used as starting point for the next iteration until the algorithm converges.

The values of the resulting estimates $\Theta^{[T]}$ do not necessarily correspond to points in the dataset. If this is the case, they are replaced by their closest point in $\mathcal{D}$ to ensure that the resulting estimate is a true subset of the original dataset. Since duplicates are eliminated, this final step causes, in general, the final estimate $\Theta^{[T]}$ to contain less points than in the original estimate $\Theta^{[0]}$.

The time complexity of the fuzzy k-means algorithm is $O(TKN)$, where T is the number of iterations required for convergence, N the number of samples in $\mathcal{D}$, and K the number of components in the model. The likelihood of the final model and the number of iterations T needed for convergence depend strongly on the distribution of the samples in the initial model $\Theta^{[0]}$. If the initial estimates are very good, convergence can be quite rapid. However, like all hill-climbing procedures, the algorithm only finds local maxima. There is no guarantee that the solution found constitutes a global optimum.

IV. GENERATING THE INITIAL MAP

The maps produced by the algorithm described in the previous section depend strongly on the initial model $\Theta^{[0]}$. A common technique for generating an initial model is to randomly select points out of the dataset $\mathcal{D}$. However, this strategy produces maps that, in general, do not represent the structure of the environment thoroughly and are, in general, not well suited for navigation tasks. In this section we present grid-based sampling as an alternative algorithm for generating an initial model. Grid-based sampling tries to spread the points along the underlying structure of the environment to get a more even distribution of the samples.

Fig. 2. The map on the left was generated by randomly selecting samples from the dataset. Observe how the samples are unevenly distributed, concentrated in areas where the original distribution has a high sample density. The map on the right was generated using grid-based sampling and contains the same number of points as the map on the left.

Figure 2 depicts two maps of 389 points each. The map on the left was generated by randomly selecting points from the original dataset. Observe how the points are unevenly distributed, concentrated in areas where the original distribution has a high sample density. The map on the right was generated using grid-based sampling. Observe how the points present a more even distribution, representing the structure of the environment more thoroughly.

A. Grid-based Sampling

Grid-based sampling takes as input a set $\mathcal{D}$ of points and returns a subset $\mathcal{S} \subseteq \mathcal{D}$. The algorithm divides the environment into cells of equal size and replaces the points that fall within the same cell with a common representative point. The resulting set $\mathcal{S}$ of samples is constituted by the mean centers m_{c_i} [17] of each non-empty cell C_i. The mean center m_{c_i} of a cell C_i is a point $m_{c_i} \in C_i$ such that

$$\sum_{x \in C_i} ||m_{c_i} - x||^2 \leq \sum_{x \in C_i} ||z - x||^2 \qquad \forall z \in C_i \,. \qquad (6)$$

The number of samples in the resulting set depends on the resolution of the grid. A finer resolution will produce a larger set than a coarser one.

The time complexity for building the grid is $O(N)$, where N is the size of $\mathcal{D}$. The mean center of a cell can be computed in $O(n)$, where n is the number of points in the cell, by keeping track of the mean of the points in the cell as the grid is build. At the end, we just need to look within the cell for the closest sample to the mean, what can be done in $O(n)$ using a naive search. The total complexity of the algorithm is given by $O(N + Kn)$, where K is the number of non-empty cells in the grid and n is the average number of points per non-empty cell.

Grid-based sampling does not attempt to find an optimal model according to a global evaluation criterion, but utilizes an approximative heuristic. It tries to partition the points according to regions, assuming that each region describes an equally important part of the environment. The main disadvantage of grid-based sampling is that the exact number of resulting samples can not be specified directly.
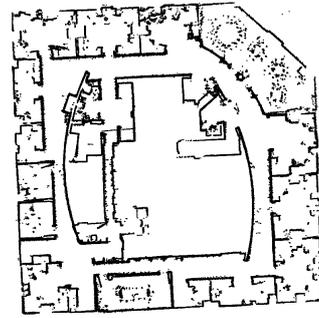
Fig. 3. The Intel Research Lab in Seattle (29m×29m) with 155,648 samples.

V. EXPERIMENTAL RESULTS

The algorithms presented above were implemented and evaluated using various datasets gathered with real robots. The goal of the experiments presented in this section is to demonstrate how our approach can be used to generate maximum-likelihood sample-based maps of the environment from laser range measurements. We also demonstrate how the improved sample-based maps increase the accuracy during robot localization when compared to the maps before improvement.

The datasets used consisted of different structured and unstructured environments that differ considerably in the number and distribution of the samples. In particular, the results presented in this section correspond to the Intel Research Lab dataset freely available online from *The Robotics Data Set Repository* (Radish) [9]. When projecting the aligned range measurements, only readings shorter than 10 meters were considered. The resulting dataset is conformed by 155,648 samples (see Figure 3). Similar results as the ones presented in this section were obtained when using other datasets.

A. Likelihood Maximization

The first experiment is designed to illustrate how our approach can be used to obtain sample-based maps that maximize the likelihood of range data gathered by the robot. We tested our approach using maps of different sizes. To initialize the algorithm grid-based sampling (Section IV) was used. We also compared our results using random sampling to generate the initial map. Figure 4 plots the likelihood for the resulting maps using different map sizes. The size of the maps is expressed as the fraction of samples in the whole dataset. The resulting maps are denoted with *+f-km* and the different initializations as *gbs* and *rnd* for grid-based sampling and random sampling respectively.

As can be seen in the figure, our approach does increase the likelihood for the maps that were used to initialize the algorithm. A two-sample t test revealed that the improvement in the likelihood is significant on the $\alpha = 0.01$ level for all the evaluated sizes and initialization techniques. The amount of improvement depends on the distribution and number of points the initial map, and as the number of points in the map increases, the amount of improvement that can be obtained through our approach decreases. In proportionally large maps, the contribution of each point to the likelihood of the data
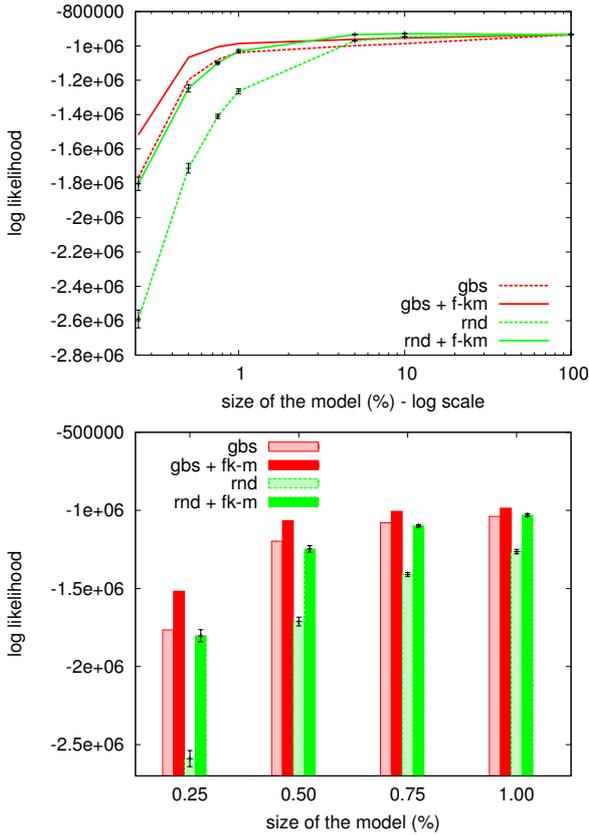
Fig. 4. Likelihood for maps of different sizes obtained using different initialization techniques: grid-based sampling (*gbs*) and random sampling (*rnd*). The corresponding resulting maps are denoted with (+*f-km*).

Fig. 5. Average localization error $d(x, x')$ (7) during position tracking for maps of different sizes obtained using different initialization techniques: grid-based sampling (*gbs*) and random sampling (*rnd*). The corresponding resulting maps are denoted with (+*f-km*).

is relatively small. Thus, as the number of points in the map increases, the strategy for selecting the samples becomes less important, since all points are almost equally relevant. The top plot in Figure 4 reveals that for maps containing more than 10% of the points of the dataset, the sub-sampling strategy becomes essentially unimportant. For proportionally large models, the application of our algorithm is no longer justified. However, we are only interested in maps that are considerably smaller than the original dataset, where the application of our algorithm provides an important increment in the likelihood. The proportional size of the map at which the sub-sampling strategy to used becomes unimportant depends on the specific dataset. In particular, it depends on the sample density of the complete dataset. The graph at the bottom of Figure 4 plots the likelihood for the resulting maps for sizes smaller than 10% of the samples of the complete dataset.

The influence of the initial map on the result of our algorithm can be appreciated in Figure 4. Observe how when randomly initializing the algorithm, the likelihood of the resulting map is sometimes smaller, than the likelihood of a map of the same size obtained using grid-based sampling even without applying our algorithm. In general, different initialization techniques produce different results. A standard approach to alleviate this problem is to run the algorithm several times using different initializations and to keep the best result.

B. Localization Using Sample-Based Maps

This experiment is designed to evaluate sample-based maps generated using our approach in the context of position tracking. We compare the accuracy of the estimated pose of the robot using Monte Carlo localization [6]. We use the maps obtained in the previous experiment and compare the accuracy of the localization algorithm in both the initial and improved maps. As error between the true and the estimated pose we use the distance $d(x, x')$ described in [14] as

$$d(x, x') = \left[\xi \delta(x_3 - x'_3)^2 + \right.$$
$$\left. (1 - \xi)((x_1 - x'_1)^2 + (x_2 - x'_2)^2) \right]^{1/2}, \qquad (7)$$

where x_1 and x'_1 are the x-coordinates, x_2 and x'_2 are the y-coordinates, and $\delta(x_3 - x'_3)$ is the difference in the orientation of x and x'. Additionally, ξ is a weighing factor that was set to 0.8 in all our experiments. As ground truth we use the corrected pose for the measurements obtained using the scan-matching algorithm of the Carnegie Mellon Robot Navigation Toolkit (CARMEN) [15]. In order to make our results general, we add some Gaussian noise to the odometry of the robot. As observation model for the localization we use the likelihood fields model as described in [19].

Figure 5 plots the average localization error obtained when using the different maps from the previous experiment. For the

particle filter implementation of the Monte Carlo localization, 500 particles were used and the average of the particles was used as estimation of the current pose of the robot. The localization algorithm was run multiple times for each map and the errors were averaged. For the selected map sizes there were no serious localization errors from which the particle filter could not recover from. When using less than 0.25% of the samples, the particle filter did diverge in some runs. It can be observed in the topmost plot of Figure 5 that using more than 1% of the samples did not always improve the accuracy of the localization, and in some cases even made it slightly worst. The reason for this lies in the used observation model. The likelihood fields model neglects the data association problem. Adding more samples to the map increases the likelihood associated to slightly delocalized particles, reducing the overall localization accuracy. The size of the map for which accurate localization could no longer be guaranteed, or the size for which additional samples cease to improve the localization accuracy depend on the specific environment.

As Figure 5 suggests, the improved maps obtained using our approach are better suited for localization tasks. The localization error is in general smaller when using the improved representations. This is specially true for small maps, where the choice of the samples is more relevant. It can also be observed that small maps that try to distribute the samples uniformly throughout the structure, like the ones obtained using grid-based sampling (*gbs*) are more appropriate for localization when compared with maps of the same size that only concentrate the samples in some parts of the environment, like random sampling (*rnd*).

VI. CONCLUSIONS

In this paper, we present a novel approach for learning point-based maps that maximize the likelihood of the sensor measurements. The key idea of our approach is to obtain a non-parametric representation which uses a subset of the data as its best model. We consider this representation as a probabilistic model and apply a variant of the fuzzy k-means algorithm to find a map that maximizes the likelihood of the whole data.

Our approach was evaluated using different datasets gathered with real robots. Experimental results illustrate how our approach can be used to generate maximum-likelihood sample-based maps of the environment from laser range measurements. Results also indicate that the improved sample-based maps increase the accuracy during robot localization when compared to the maps before improvement. We believe that maximum-likelihood sample-maps constitute an accurate and general representation of the environment and are therefore suitable for robotic applications.

Despite these encouraging results, the presented approach suffers from several limitations that are the aim for future research. First, our approach is strongly dependent on the map used for initialization, and gets easily stock in local maxima. In the future we therefore will investigate strategies to scape local maxima. Another limitation is that the number of samples in the resulting maps is fixed and must be specified beforehand. This can be justified if the user wants to fix the size of the resulting maps. But in general, we would like to determine the size of the resulting maps while generating it. Another interesting area for future work would be to try our approach beyond two dimensions. The framework can be easily extended to three dimension, for example, an evaluated with three-dimensional maps.

REFERENCES

[1] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. David Levin, and C. T. Silva. Point set surfaces. In *Proc. of the conference on Visualization (VIS)*, 2001.
[2] Kai O. Arras and Roland Siegwart. Feature extraction and scene interpretation for map-based navigation and map building. In *Proc. of SPIE, vol. 3210, Mobile Robotics XII, p. 42-53*, Pittsburgh, USA, 1997.
[3] Peter Biber and Tom Duckett. Dynamic maps for long-term operation of mobile service robots. In *Robotics: Science and Systems*, pages 17–24, 2005.
[4] J. L. Crowley. World modeling and position estimation for a mobile robot using ultrasonic ranging. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1989.
[5] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley Interscience, 2000.
[6] D. Fox, W. Burgard, F. Dellaert, and S. Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proc. of the National Conference on Artificial Intelligence*, 1999.
[7] J. González, 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)*, pages 1904–1909, 1994.
[8] H. González-Baños and J.-C. Latombe. Robot navigation for automatic model construction using safe regions. In *Int. Symposium on Experimental Robotics (ISER)*, pages 405–415, 2000.
[9] A. Howard and N. Roy. The robotics data set repository (radish), http://radish.sourceforge.net/, 2003.
[10] J. Leonard, P Newman, R. Rikoski, J. Neira, and J. Tardós. Towards robust data association and feature modeling for concurrent mapping and localization. In *10th Internation Symposium on Robotics Research (ISRR)*, 2001.
[11] M Levoy and T. Whitted. The use of points as a display primitive. Computer Science Technical Report 85-022, UNC-Chapel Hill, 1985.
[12] H. P. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1985.
[13] M. Pauly, M. Gross, and L. P. Kobbelt. Efficient simplification of point-sampled surfaces. *Theor. Comput. Sci.*, 284(1):67–108, 2002.
[14] P. Pfaff, W. Burgard, and D. Fox. Robust monte-carlo localization using adaptive likelihood models. In H.I. Christiensen, editor, *European Robotics Symposium 2006*, volume 22, pages 181–194. Springer-Verlag Berlin Heidelberg, Germany, 2006.
[15] N. Roy, M. Montemerlo, and S. Thrun. Perspectives on standardization in mobile robot programming. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
[16] D. Schröter, M. Beetz, and J.-S. Gutmann. Rg mapping: Learning compact and structured 2d line maps of indoor environments. In *Proceedings of the IEEE International Workshop on Robot and Human Interactive Communication (ROMAN'02)*, 2002.
[17] S. Theodoridis and K. Koutroumbas. *Pattern Recognition*. Elsevier Academic Press, 2003.
[18] S. Thrun. Learning occupancy grids with forward sensor models. In *Autonomous Robots*, volume 15, pages 111–127, 2003.
[19] S. Thrun, D. Fox, W. Burgard, and Dellaert. F. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2), 2001.
[20] M. Veeck and W. Burgard. Learning polyline maps from range scan data acquired with mobile robots. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.