

Using artificial landmarks to reduce the ambiguity in the environment of a mobile robot

D. Meyer-Delius, M. Beinhofer, Alexander Kleiner and W. Burgard

Post Print

N.B.: When citing this work, cite the original article.

©2011 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

D. Meyer-Delius, M. Beinhofer, Alexander Kleiner and W. Burgard, Using artificial landmarks to reduce the ambiguity in the environment of a mobile robot, 2011, IEEE Int. Conf. on Robotics and Automation (ICRA), 5173-5178.

<http://dx.doi.org/10.1109/ICRA.2011.5980111>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-72521>

Using Artificial Landmarks to Reduce the Ambiguity in the Environment of a Mobile Robot

Daniel Meyer-Delius

Maximilian Beinhofer

Alexander Kleiner

Wolfram Burgard

Abstract—Robust and reliable localization is a fundamental prerequisite for many applications of mobile robots. Although there exist many solutions to the localization problem, structurally symmetrical or featureless environments can prevent different locations from being distinguishable given the data obtained with the robot’s sensors. Such ambiguities typically make localization approaches more likely to fail. In this paper, we investigate how artificial landmarks can be utilized to reduce the ambiguity in the environment. We present a practical approach to compute a configuration of indistinguishable landmarks that decreases the overall ambiguity and thus increases the robustness of the localization process. We evaluate our approach in different environments based on real data and in simulation. Our results demonstrate that our approach improves the localization performance of the robot and outperforms other landmark selection approaches.

I. INTRODUCTION

For reliable navigation a mobile robot needs to be able to determine its pose in the environment and accurately track it over time. This is known as the mobile robot localization problem and consists in estimating the pose of the robot relative to a given map of the environment based on sensor data. Although there exist many approaches that have been successfully applied to the localization task, structurally symmetrical or featureless environments make these approaches more likely to fail and in the worst case prevent the pose of the robot from being uniquely determinable at all. In the context of localization, environments are considered ambiguous if they prevent different robot poses from being distinguished based on the sensor data. Figure 1 illustrates such a problem. It shows a typical sensor measurement obtained using a laser range scanner together with the corresponding observation likelihood. Dark colored areas correspond to high likelihood poses. As can be seen in the figure, several poses, in addition to the pose from which the scan was taken, have a high observation likelihood.

In this paper, we focus on the problem of utilizing artificial landmarks to reduce the ambiguity in the environment. Concretely, we address the problem of finding a configuration of indistinguishable landmarks that, when placed in the environment, increase the robustness in the localization of the robot. The basic idea of our approach is that by reducing the overall ambiguity in the environment, the localization performance of the robot can be improved. We first introduce




Fig. 1. Ambiguous environments make localization approaches more likely to fail. The image shows a laser range scan (top) and the x - y -projection (maximizing over the orientation) of the corresponding observation likelihood over the pose space (bottom). The darker the color, the larger the observation likelihood at the corresponding pose. As can be seen there are many poses with an associated high observation likelihood.

a measure for how distinguishable or unique a pose is. We then present a landmark selection approach that incrementally selects landmark locations, greedily maximizing the average uniqueness in the environment.

The contribution of this paper is two-fold: First, we present a practical approach to landmark placement that aims at improving the localization performance of the robot. Our approach provides us with both the number and location of landmarks to be placed in the environment. We consider indistinguishable landmarks which makes our approach attractive from a practical point of view since no landmark coding and complex identification system is required. As a second contribution we introduce a measure for the uniqueness of a robot pose based on sensor data. In this work, we seek to improve the localization performance by maximizing the average uniqueness in the environment. Furthermore, we describe a concrete instantiation of the landmark placement problem and show experimentally that our approach improves the localization performance of the robot and outperforms other landmark selection approaches.

II. RELATED WORK

In the past, a huge variety of localization techniques has been proposed. Several approaches rely on natural features in the environment [1], [2]. These approaches are particularly attractive as they do not require the environment to be modified. However, inherently ambiguous environments make these localization approaches more likely to fail. Utilizing artificial landmarks offers the possibility of improving the reliability of the localization. Some approaches [3], [4] consider artificial landmarks that can be uniquely identified. Although such approaches greatly simplify the localization

All authors are with the University of Freiburg, Department of Computer Science, D-79110 Freiburg, Germany.

This work has been partially supported by the European Commission under contract numbers FP7-248258-First-MM and FP7-260026-TAPAS. Also by the German Research Foundation (DFG) under contract number SFB/TR 8 and within the Research Training Group 1103.

problem, they require a landmark coding and non-trivial identification system. In our work, we therefore consider only indistinguishable landmarks.

The landmark placement problem as addressed in our work can be formulated as the problem of selecting a subset of landmarks out of a finite set of candidate landmarks. Sutherland and Thompson [5] were one of the first to address this problem. They demonstrate that the localization error depends on the configuration of the selected landmarks. Salas and Gordillo [6] propose a simulated annealing technique to find the landmark configuration that maximizes the size of the region from where a landmark can be seen. Sinriech and Shoval [7] specify a set of constraints about the number of landmarks and their distance to critical locations in the environment, and formulate landmark placement as a nonlinear optimization problem. Sala *et al.* [8] decompose the environment into regions from which a minimum number of landmarks can be observed. They use a graph-theoretical formulation to find the decomposition with the minimum number of regions. All of the above mentioned approaches rely on pure geometrical reasoning based on triangulation for estimating the pose of the robot. In contrast to that, our approach to landmark selection is tightly coupled with a robust, probabilistic localization framework.

Other researchers have also focused on the localization performance at the moment of selecting landmarks. Thrun [9], for example, uses a neural network to extract features from the sensor data and selects the subset of those features that minimizes the average posterior localization error. Lerner *et al.* [10] formulate the problem as a semi-definite programming (SDP) problem and specify a cost function to weight different localization parameters according to the specific task at hand. Strasdat *et al.* [11] use reinforcement learning to obtain an online landmark selection policy. The approach of Zhang *et al.* [12] selects, at every time step, the set of landmarks that minimizes the entropy of the resulting posterior distribution. All of these methods operate online and are concerned with the landmarks that are observed at every time step during localization. Our approach, in contrast, works in an offline fashion. Furthermore, in our work landmarks are physical objects, and not observed features. The main difference between previous approaches and ours is that, by introducing a measure for the uniqueness of a pose and maximizing the overall uniqueness in the environment, we explicitly consider the symmetries and ambiguities that can originate when placing landmarks.

III. MONTE CARLO LOCALIZATION

Throughout this work we use the Monte Carlo localization (MCL) [13] algorithm for estimating the pose of the robot. MCL is a concrete instantiation of the Bayes recursive filter that estimates the belief $bel(x_t) = p(x_t | z_{1:t}, u_{0:t}, m)$ about the pose x_t of the robot at time t conditioned on the observations $z_{1:t}$, odometry measurements $u_{0:t}$, and map of the environment m . The belief is recursively computed as

$$\overline{bel}(x_t) = \int p(x_t | x_{t-1}, u_t) bel(x_{t-1}) dx_{t-1} \quad (1)$$

$$bel(x_t) = \eta p(z_t | x_t, m) \overline{bel}(x_t), \quad (2)$$

where $p(z_t | x_t, m)$ and $p(x_t | x_{t-1}, u_t)$ correspond to the observation and motion model respectively, and η is a normalizing constant independent of x_t . Equation (1) is known as the prediction step of the algorithm and computes the posterior $\overline{bel}(x_t)$ before incorporating the observation z_t . Equation (2) is called the correction step and computes the final belief $bel(x_t)$. The key idea of MCL is to represent the belief by a set of weighted samples or particles, where each particle corresponds to a potential pose of the robot.

The MCL algorithm computes the particle set at time t recursively from the particle set at time $t-1$. The algorithm is initialized with a distribution of equally weighted particles around the initial pose estimate (position tracking), or with a uniform distribution over all possible poses (global localization). Then, a temporary particle set is generated from the previous set by sampling according to the motion model $p(x_t | x_{t-1}, u_t)$. This model describes a posterior density over possible poses x_t given the previous pose x_{t-1} and most recent odometry measurement u_t . In a following step, the new particles are weighted according to the observation model $p(z_t | x_t, m)$ that represents the likelihood of the most recent observation z_t given the pose x_t and the map m of the environment. In the final step of the algorithm, the resulting particle set is created by drawing particles from the temporary set in proportion to their weights. This resampling step effectively replaces particles with low weights by high weight particles.

IV. POSE UNIQUENESS

Intuitively, the uniqueness of a pose indicates how distinguishable the pose is from all other poses in the state space. Since the robot perceives its environment through its sensors, the uniqueness of a pose is based on the observations of the robot. Let us assume that the robot is equipped with a perfect sensor that makes, for a pose x and map m , a deterministic observation $z(x, m)$. Then we can define the *uniqueness* of a pose x given a map m as

$$\mathcal{U}_{\text{perfect}}(x, m) = \frac{1}{\int_{\tilde{x} \in \mathcal{X}} \delta_{\tilde{x}, m}(x) d\tilde{x}}, \quad (3)$$

where \mathcal{X} represents the state space, and $\delta_{\tilde{x}, m}(x)$ is 1 if $z(\tilde{x}, m) = z(x, m)$ and 0 otherwise. The denominator in (3) simply counts the number of poses in the state space where the robot makes the same observation as in pose x . Clearly, the larger the count, the less unique the pose is. For a maximally unique pose x it holds that $\delta_{\tilde{x}, m}(x) = 0$ for all $\tilde{x} \in \mathcal{X} \setminus \{x\}$. A minimally unique pose x , on the other hand, is one for which $\delta_{\tilde{x}, m}(x) = 1$ for all $\tilde{x} \in \mathcal{X}$.

Since our sensor is noisy, we have to replace the deterministic function $\delta_{\tilde{x}, m}(x)$ in (3) by the likelihood of observing, at pose \tilde{x} , the observation z^x made at x , i.e., $p(z^x | \tilde{x}, m)$. Furthermore, as we don't know which measurement z^x we will obtain at pose x , we have to integrate over all potential measurements, i.e., calculate the expectation:

$$\mathcal{U}_{\text{exp}}(x, m) = \int_z \frac{1}{\int_{\tilde{x} \in \mathcal{X}} p(z | \tilde{x}, m) d\tilde{x}} p(z | x, m) dz. \quad (4)$$

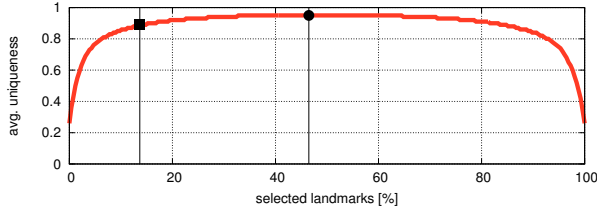


Fig. 2. Normalized average uniqueness as a function of the number of selected landmarks. The right black dot marks the point where the average uniqueness reaches its maximum. The left black square indicates the average uniqueness for the number of landmarks selected by our approach.

As integrating over all measurements is not feasible in practice, we approximate the outer integral in $\mathcal{U}_{\text{exp}}(x, m)$ by the maximum likelihood value of $p(z | x, m)$ and obtain

$$\mathcal{U}(x, m) = \frac{1}{\int_{\tilde{x} \in \mathcal{X}} p(z^{x^*} | \tilde{x}, m) d\tilde{x}}, \quad (5)$$

where $z^{x^*} = \operatorname{argmax}_z p(z | x, m)$ corresponds to the most likely observation at pose x given the map m . The accuracy of this approximation depends on the distribution $p(z | x, m)$. If $p(z | x, m)$ is the Dirac density, the approximation is exact. In general, the accuracy depends on how much z^{x^*} dominates the outer integral. A highly unique pose x is typically associated to a peaked distribution $p(z | \tilde{x}, m)$ that is at its maximum when $\tilde{x} = x$. On the other hand, a flat distribution typically corresponds to an ambiguous environment where, for all the poses \tilde{x} in the state space, $p(z | \tilde{x}, m)$ has almost the same value.

V. LANDMARK PLACEMENT

Given a set \mathcal{V} of N candidate landmarks, the general landmark placement problem consists of finding a configuration $m \subseteq \mathcal{V}$ of landmarks that maximizes a given target function. There exist many possible aspects to consider when specifying the target function, like the number of selected landmarks and area covered, for example. The target function considered in this work is the average uniqueness value in the environment. Concretely, we look for the configuration m^* so that

$$m^* = \operatorname{argmax}_{m \subseteq \mathcal{V}} \left(\frac{1}{|\mathcal{X}|} \int_{x \in \mathcal{X}} \mathcal{U}(x, m) dx \right). \quad (6)$$

By maximizing the average uniqueness in the environment we seek to improve the localization performance of the robot.

The combinatorial nature of the problem makes the enumeration of all possible solutions for finding the optimal one intractable. However, an approximate solution to (6) can be efficiently computed in an incremental fashion by successively selecting the landmark that maximizes the average uniqueness until no further improvement is possible. The main disadvantage of this approach is that it selects an unnecessarily large number of landmarks. In practical experiments we found that approximately 50% of the candidate landmarks are selected before no further improvement is possible. Figure 2 shows the typical behavior of the average uniqueness as a function of the number of selected landmarks. As can be seen, the average uniqueness reaches

Algorithm 1 Incremental Landmark Placement

Require: Set \mathcal{V} of N candidate landmarks

```

1:  $m^* = \emptyset$ 
2: while  $\mathcal{V} \neq \emptyset$  do
3:    $l' = \operatorname{argmax}_l \overline{\mathcal{U}}(x, \{l\} \cup m^*)$ 
4:    $\bar{u}' = \overline{\mathcal{U}}(x, \{l'\} \cup m^*)$ 
5:   if  $\nabla \bar{u}' > \epsilon$  then
6:      $m^* = \{l'\} \cup m^*$ 
7:      $\mathcal{V} = \mathcal{V} \setminus \{l'\}$ 
8:   else
9:     return  $m^*$ 

```

its maximum when approximately half of the candidate landmarks are selected. Adding further landmarks provides no additional improvement and, as a matter of fact, the average uniqueness starts to decrease as further landmarks are selected. In order to determine the number of landmarks to select, we use a heuristic approach based on the gradient of the average uniqueness. Our landmark placement algorithm terminates whenever the gradient drops below some specified threshold. The larger the value for the threshold, the smaller the number of selected landmarks. The average uniqueness is normalized using an instance specific upper bound in order to use the same threshold for different instances of the problem. Assuming a finite and discrete state space, the upper bound for the average uniqueness is given by $1/\min_x p(z^{x^*} | x, m = \emptyset)$ and can be determined before selecting the first landmark.

The approach proposed in this work is specified in Algorithm 1. Line 3 computes the landmark l' that maximizes the average uniqueness $\overline{\mathcal{U}}(x, \{l'\} \cup m^*)$. Line 5 computes the gradient of the average uniqueness if l' would be selected. If the value of the gradient is smaller than the threshold ϵ , the algorithm terminates and the final configuration m^* is returned. Otherwise, landmark l' is added to the configuration m^* , removed from the set of candidate landmarks, and the algorithm continues and tries to select another landmark. Assuming a finite and discrete state space, the complexity of the algorithm is $O(KM^2)$ where $K < N$ is the number of selected landmarks and M is the size of the state space. The factor M^2 is a consequence of the computation of the average uniqueness that requires the computation of the uniqueness $O(M)$ for each state in the state space.

VI. EXPERIMENTAL EVALUATION

To evaluate the improvement in the localization performance obtained when landmarks were placed according to our proposed algorithm, we carried out a set of experiments in simulation and on a real robot. A 2-dimensional occupancy grid with a resolution of $0.5m$ was used to represent the environment. The set of candidate landmarks consisted of all occupied cells in the grid. The sensor used for our experiments was a laser range scanner that in addition to the range and bearing, also returned the reflectivity of the measured objects. As landmarks we considered stripes of retro-reflective tape. Based on the reflectivity we classified




Fig. 3. Improving the uniqueness in the environment by placing landmarks. The figure shows the uniqueness before and after placing the landmarks for three different environments. The uniqueness at each pose is projected onto the grid map by minimizing over the orientation. The lighter the color, the higher the uniqueness. Also shown are the landmark configurations obtained using our approach. The highlighted locations correspond to the landmarks.

individual measurements into pure range measurements and measurements that correspond to landmarks.

As observation model $p(z | x, m)$ we used a variant of the so-called likelihood field model [14]. In this model, the individual range measurements are assumed to be independent of each other. The likelihood of each pure range measurement is computed according to the distribution $p_r(z^i) \sim \mathcal{N}(d_r, \sigma_r^2)$ based on the distance d_r between the endpoint of the measurement z^i and its closest obstacle in the map m . Here, $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and standard deviation σ . In the case of landmark measurements the distribution $p_l(z^i) \sim \mathcal{N}(d_l, \sigma_l^2)$ is used, where d_l denotes the distance to the closest landmark in the map. The likelihood of an observation $z = (z^1, \dots, z^K)$ is then computed as

$$p(z | x, m) = \prod_{i=1}^K p_l(z^i)^{\delta(z^i)} \cdot p_r(z^i)^{(1-\delta(z^i))}, \quad (7)$$

where

$$\delta(z^i) = \begin{cases} 1 & \text{if } z^i \text{ corresponds to a landmark} \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

This simple general model does not take visibility constraints into account and assumes a perfect landmark detection. However, it can be efficiently evaluated and is sufficient for the purpose of our experiments.

To compute the uniqueness as specified in (5), the state space (x - y coordinates and orientation θ) was divided into cells of $0.5m$ and a resolution of 90° was used for the orientation. Ray-tracing was used to simulate the most likely observations needed to compute the uniqueness value at every pose. Figure 3 shows the landmark configurations obtained with our approach for three artificial maps. Also shown in the maps is the uniqueness in x - y space, minimized over the orientation θ , before and after placing the landmarks. As threshold ϵ in our landmark placement algorithm we set the minimum value of the gradient to 1. For our specific sensor model (7), a theoretical upper bound for the average uniqueness is given by $1/\sqrt[3]{2\pi\sigma^2}$, where $\sigma = \max(\sigma_r, \sigma_l)$. Note that the approach presented in this paper is not restricted to grid-based representations, it only requires a way

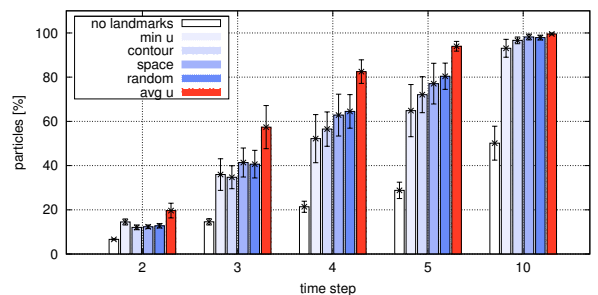


Fig. 4. Evaluation of the landmark configuration obtained using our approach vs. alternative approaches in the task of global localization. The plots show the fraction of particles within a $1.5m$ radius around the true pose as a function of time. The values correspond to the mean and standard deviation for different repetitions of the localization.

to compute the uniqueness for the specific representation. Additionally, the computation of the upper bound for the average uniqueness is not strictly necessary. Alternatively, the empirical maximal value can be used instead. The disadvantage of that strategy is that the algorithm can't stop until the maximal average uniqueness value has been reached.

A. Global Localization

To evaluate the landmark configurations obtained using our approach in the task of global localization, we generated 50 different random trajectories for each of the environments shown in Figure 3. In addition to noise in the range simulations, we also simulated false positives and false negatives in the landmark detections. The localization algorithm was executed 5 times for each trajectory using 10 000 particles initially uniformly distributed in the state space. We compared, for the same number of landmarks, the configurations obtained with our approach (*avg u*) with the configurations obtained with four alternative approaches:

- 1) *Uniform contour sampling (contour)*: Distributes the landmarks roughly uniformly throughout the contours of the environment. The first landmark is randomly selected from the set of candidate landmarks. Additional landmarks are selected by choosing the candidate landmark closest to the previously selected one until no more landmarks can be selected. Every time a

landmark is selected, all landmarks within a specified radius are removed from the set of candidates. This radius is chosen so that the sampling approximately covers the whole map.

- 2) *Uniform space sampling (space)*: Distributes the landmarks roughly uniformly throughout the environment. This approach divides the environment into squared regions of equal size and selects the candidate landmark closest to the center of each non-empty region.
- 3) *Random sampling (random)*: Distributes the landmarks randomly throughout the contours of the environment. Landmarks are randomly selected from the set of candidate landmarks. Every time a landmark is selected, all landmarks within a specified radius are removed from the set of candidates.
- 4) *Maximize the minimal uniqueness (min u)*: The approach described in Algorithm 1 was modified so that it would maximize the minimum uniqueness in the environment $m^* = \operatorname{argmax}_{m \subseteq \mathcal{V}} (\min \mathcal{U}(x, m))$, instead of maximizing the average uniqueness,

The results of the experiment for one of the environments (leftmost in Figure 3) are shown in Figure 4. As performance metric for the global localization task we considered the fraction of particles within a 1.5 m radius around the true pose after 2, 3, 4, 5 and 10 integrations of measurements (time steps). The values correspond to the mean and standard deviation for the different trajectories and runs of the localization. As can be seen in the figure, the configuration obtained by our method improves the global localization performance best since particles are more quickly converging towards the true pose of the robot. A t-test showed that the improvement was significant on the $\alpha = 0.05$ level for all the evaluated environments, time steps and alternative approaches. Clearly, the amount of improvement obtainable depends on the inherent uniqueness of the environment. A larger improvement can be obtained for inherently ambiguous environments (leftmost one in Figure 3) than for inherently unique environments (rightmost one in Figure 3). The first 3 alternative approaches, *contour*, *space*, and *random* are simple and fast, but do not take into account the ambiguities that can originate when selecting landmarks, and the resulting landmark configurations are therefore not as good for improving the localization performance as the ones obtained with our approach. The fourth approach, *min u*, has the property that a lower bound for the uniqueness in the environment is guaranteed. However, this does not provide a significant improvement in the localization performance.

B. Number of Landmarks

The goal of this experiment was to evaluate the performance of our gradient-based heuristic when determining automatically the number of landmarks to be placed. Figure 5 shows the fraction of particles within a 1.5 m radius around the true pose of the robot after 2, 3, and 10 integrations of measurements as a function of the number of selected landmarks. Also indicated in the figure are the

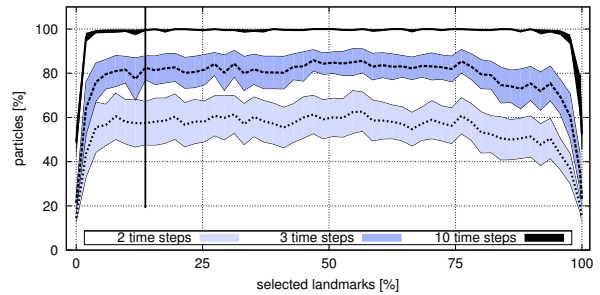


Fig. 5. Global localization performance as a function of the number of selected landmarks. The plots show the fraction of particles within a 1.5 m radius around the true pose of the robot after 2, 3, and 10 time steps. The values correspond to the mean and standard deviation for different repetitions of the experiment. The number of landmarks selected by our approach ($\sim 10\%$) is indicated by the vertical line.

values corresponding to the number of landmarks selected by our gradient-based heuristic ($\sim 10\%$) when using a threshold of 1. This corresponds to a 45° positive gradient. The motivation for choosing this value is that increasing the percentage of selected landmarks by 1% provides less than a 1% increment in the normalized average uniqueness. Using a different value for the threshold, or weighting differently the parameters (average uniqueness vs. fraction of selected landmarks) the number of selected landmarks can be controlled. As can be seen in Figure 5 selecting more landmarks does not provide an improvement in the localization performance. Furthermore, the number of landmarks selected is well beyond the point where fewer landmarks would cause the localization performance to decrease drastically.

An additional result of this experiment is that, as can be seen in the figure, the behavior of the localization performance as a function of the number of selected landmarks is similar to the behavior of the average uniqueness (see Figure 2). This experimental result suggests a direct connection between the average uniqueness in the environment and localization performance.

C. Real Data

We also evaluated our approach using data gathered with a MobileRobots Pioneer P3-DX robot equipped with a SICK LMS 291 laser range finder. We steered the robot through one of the buildings at our campus and created an occupancy map of the environment using a standard SLAM technique [15]. The environment consisted of a long, featureless corridor of approximately $80 \times 3\text{ m}$ size. Figure 6 shows the landmark configuration obtained using our approach. In order to make the environment more ambiguous, range measurements larger than 10 m were disregarded. For building the map, however, the full 80 m depth range of the laser scanner was used. The map resolution and state space discretization are described in Section VI. As landmarks, stripes of retro-reflective material were taped to the walls in the locations indicated by our approach. We used a threshold on the reflectivity value to classify the laser measurements caused by the landmarks.

After placing the landmarks, we steered the robot again through the environment and used the above mentioned



Fig. 6. Landmark configuration and uniqueness before and after placing the landmarks for the *building 103* data set gathered with a real robot using a laser range scanner.




Fig. 7. Global localization performance for the *building 103* data set. The plots show the fraction of particles within a 1.5 m radius of the true pose of the robot as a function of time. Also shown are the results obtained in simulation (*sim*).

SLAM technique to obtain an approximated ground truth for comparison. For the statistical analysis we divided the data into 5 parts and evaluated the landmark configuration in the task of global localization as described in Section VI-A. Figure 7 shows the results of the experiment. As expected, a substantial improvement in the localization performance was obtained when using the landmarks. The performance is, however, lower than the one obtained in simulation. This is mostly due to the simplistic sensor model considered for the experiments. The model is sufficient for the purpose of our evaluation, but we expect that using a better model, for example one that considers the distance and angle of incidence of the beams at the moment of detecting a landmark, would produce better results for real data.

We also evaluated the configuration obtained using our approach in the task of position tracking. Figure 8 shows the Euclidean distance between the average pose of the particle set and the true pose of the robot as a function of time. As can be seen in the figure, using landmarks can also improve the accuracy of the localization in the ambiguous areas of the environment. With and without landmarks the error grows as the robot moves along the corridor and decreases when the ends of the corridor are visible ($t \sim 150, 300, 450, 600$). With landmarks, however, the error is substantially smaller when moving along the corridor.

VII. CONCLUSIONS

In this paper we presented a landmark placement approach that seeks to reduce the overall ambiguity in the environment to improve the localization performance of a mobile robot. To this extend we proposed a measure for the uniqueness of a robot pose based on the appearance of the environment as observed by the robot. Due to the combinatorial nature of the landmark placement problem, we introduced an approximative approach that incrementally selects landmark locations




Fig. 8. Average localization error during position tracking with and without utilizing landmarks for the *building 103* data set.

from a set of candidate locations and thereby maximizes the average uniqueness in the environment. Furthermore, we described a concrete application in the context of localization with laser range scanners given a grid-based representation of the environment. We evaluated our approach for different environments in simulation and using real data. The results demonstrate that our approach yields substantial improvements in the localization performance.

REFERENCES

- [1] J. Leonard and H. Durrant-Whyte, “Mobile robot localization by tracking geometric beacons,” *Robotics and Automation, IEEE Transactions on*, vol. 7, no. 3, pp. 376–382, jun. 1991.
- [2] W. Burgard, D. Fox, D. Hennig, and T. Schmidt, “Estimating the absolute position of a mobile robot using position probability grids,” in *Proc. of the National Conference on Artificial Intelligence*, 1996.
- [3] A. Howard, M. Mataric, and G. Sukhatme, “Relaxation on a mesh: a formalism for generalized localization,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2001.
- [4] C. Raffin and A. Fournier, “Learning with a friendly interactive robot for service tasks in hospital environments,” *Autonomous Robots*, vol. 3, pp. 399–414, 1996.
- [5] K. Sutherland and W. Thompson, “Inexact navigation,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1993.
- [6] J. Salas and J. L. Gordillo, “Placing artificial visual landmarks in a mobile robot workspace,” in *Proc. of the Ibero-American Conf. on Artificial Intelligence (IBERAMIA)*, 1998.
- [7] D. Sinriech and S. Shoval, “Landmark configuration for absolute positioning of autonomous vehicles,” *IIE Transactions*, vol. 32, pp. 613–624, 2000.
- [8] P. Sala, R. Sim, A. Shokoufandeh, and S. Dickinson, “Landmark selection for vision-based navigation,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2004.
- [9] S. Thrun, “Finding landmarks for mobile robot navigation,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1998.
- [10] R. Lerner, E. Rivlin, and I. Shimshoni, “Landmark selection for task-oriented navigation,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2006.
- [11] H. Strasdat, C. Stachniss, and W. Burgard, “Which landmark is useful? Learning selection policies for navigation in unknown environments,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2009.
- [12] S. Zhang, L. Xie, and M. Adams, “Entropy based feature selection scheme for real time simultaneous localization and map building,” in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.
- [13] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, “Monte carlo localization for mobile robots,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1999.
- [14] S. Thrun, “A probabilistic online mapping algorithm for teams of mobile robots,” *International Journal of Robotics Research*, vol. 20, no. 5, pp. 335–363, 2001.
- [15] G. Grisetti, C. Stachniss, and W. Burgard, “Improving grid-based slam with rao-blackwellized particle filters by adaptive proposals and selective resampling,” in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2005.