

Which Landmark is Useful?

Learning Selection Policies for Navigation in Unknown Environments

Hauke Strasdat

Cyryll Stachniss

Wolfram Burgard

Abstract—In general, a mobile robot that operates in unknown environments has to maintain a map and has to determine its own location given the map. This introduces significant computational and memory constraints for most autonomous systems, especially for lightweight robots such as humanoids or flying vehicles. In this paper, we present a novel approach for learning a landmark selection policy that allows a robot to discard landmarks that are not valuable for its current navigation task. This enables the robot to reduce the computational burden and to carry out its task more efficiently by maintaining only the important landmarks. Our approach applies an unscented Kalman filter for addressing the simultaneous localization and mapping problems and uses Monte-Carlo reinforcement learning to obtain the selection policy. Based on real world and simulation experiments, we show that the learned policies allow for efficient robot navigation and outperform handcrafted strategies. We furthermore demonstrate that the learned policies are not only usable in a specific scenario but can also be generalized towards environments with varying properties.

I. INTRODUCTION

In recent years, there has been a trend towards embedded systems in robotics. A series of such approaches deal with autonomous cars, helicopters, blimps, underwater vehicles, and wheeled or humanoid robots. As embedded systems typically have much higher limitations with respect to the computational power and memory capacity, it is important in the context of embedded systems to develop efficient algorithms that scale with the computational constraints of the underlying hardware.

In robotics, one of the core capabilities needed for the majority of applications is autonomous navigation. For truly autonomous navigation in initially unknown environments, the robot has to solve the so-called simultaneous localization and mapping (SLAM) problem [2, 12, 19]. Solving the SLAM problem, however, is computationally demanding and the memory requirements increase with the number of landmarks that need to be maintained by the robot. In practice, there are many scenarios in which the number of visible landmarks during a navigation task is significantly larger than the number of landmarks which can be processed efficiently using an embedded device. This leads to the question which landmark should be stored and maintained by the robot to optimally solve the navigation task. A landmark is only useful if it contributes to keep an accurate pose

estimate of the robot at the right time and in such a way that it is valuable for the navigation task. In this paper, we present an approach for learning a landmark selection policy that optimizes the navigation task carried out by the robot given its computational or memory constraints. It is obvious that the utility of a landmarks depends on the type of navigation task. We analyze two types of navigation tasks: A single-goal navigation task and a round-trip navigation task where subgoals are visited more than once. One major advantage of our approach is that the policies are not limited to the environment they have been learned in. Rather, they can also be applied successfully in environments with different properties of the underlying landmark distribution.

This paper is organized as follows. After a discussion of related work, Section III briefly introduces the unscented Kalman filter and its application to SLAM as well as reinforcement learning. Section IV then describes the different navigation tasks considered in this paper. After that, we introduce our approach to learn the optimal landmark selection policy. Finally, we present experimental results carried out in simulation as well as on a real wheeled robot.

II. RELATED WORK

The standard method for SLAM relies on the extended Kalman filter (EKF) [11] or its variants such as the unscented Kalman filter (UKF) [7]. Using these approaches, the computational requirement and memory demand increase with the number of landmarks since the full correlation between the position all landmarks is taken into account. There are many approximative filtering techniques for SLAM [12, 19]. These methods do not incorporate the full correlation between the landmarks, so that the computational constraints are less restrictive. However, their memory demand increases at least linearly with the number of landmarks used.

Recently, Sala *et al.* [15] presented a graph-theoretic formulation for the selection problem of visual features to perform navigation in known environments. The optimal set of features is defined as the minimal set with which navigation is possible. Zhang *et al.* [20] proposed an entropy-based landmark selection method for SLAM. This method specifies a measure about which visible landmark is best in the sense of entropy reduction. However, it only provides a vague guideline for how many features should be selected at a given point in time. Furthermore, Lerner *et al.* [9] presented another quality measure for landmark selection in known environments which is based on the comparison of pose uncertainties. Dissanayake *et al.* [6] suggested a map management which ensures a uniform distribution of

This work has partly been supported by the DFG within the Research Training Group 1103 and under SFB/TR-8 as well as by the European Commission under FP6-IST-34120-muFly and FP7-231888-EUROPA.

H. Strasdat, C. Stachniss, and W. Burgard are with Dep. of Computer Science, University of Freiburg, Germany. H. Strasdat is also with the Dep. of Computing, Imperial College London, UK.

landmarks over the traversed area. Apart from landmark selection, other active methods were presented such as maximizing the SLAM estimate by intelligent path planning [10] or increasing the performance of a soccer playing robot by active sensing [5].

In this paper, we present a new and universal approach for landmark selection in unknown environments. The value of a landmark is measured in terms of how well it improves the navigation/localization capabilities of the robot given the targeted navigation task. This is especially important for robots with restricted resources. We learn a landmark selection policy using Monte-Carlo reinforcement learning [3, 17, 18] and k -nearest neighbor regression [16]. We show in real world and simulation experiments that this technique allows for more efficient robot navigation.

III. PRELIMINARIES

A. Unscented Kalman Filter

The unscented Kalman filter (UKF) is a recursive Bayes' filter that estimates the state \mathbf{x} of an dynamical system in discrete time steps given a sequence of actions \mathbf{u} and observations \mathbf{z} . The n -dimensional state vector \mathbf{x} is represented by a multivariate Normal distribution with mean μ and covariance matrix $\Sigma^{(n \times n)}$. The dynamics of the system are described by a *state transition function* g plus Gaussian noise $\epsilon_{g,t}$:

$$\mathbf{x}_t = g(\mathbf{u}_t, \mathbf{x}_{t-1}) + \epsilon_{g,t} \quad (1)$$

Measurements are integrated using the *observation function*:

$$\hat{\mathbf{z}}_t = h(\mathbf{x}_t) + \epsilon_{h,t} \quad (2)$$

Again, Gaussian noise $\epsilon_{h,t}$ is added. Since Kalman filtering is an approach for systems governed by a linear difference equation, special efforts have to be taken to take the non-linearities in g and h into account.

The key idea of the unscented Kalman filter, which has been introduced by Julier and Uhlmann [7], is to apply a deterministic sampling technique that is known as the unscented transform to select a small set of so-called sigma points around the mean. Then, the sigma points are propagated through the non-linear functions. Afterwards, mean and covariance estimates are computed based on the transformed points. The advantage of this technique is that the filter can much better deal with non-linearities and thus lead to a more robust technique than the EKF.

B. Simultaneous Localization and Mapping

In the context of the SLAM problem, one seeks to simultaneously determine the map of the environment and the pose of the robot. Probabilistic methods seek to estimate the joint probability distribution

$$p(\mathbf{p}_t, \mathbf{l}_1, \dots, \mathbf{l}_M | \mathbf{u}_1, \dots, \mathbf{u}_t, \mathbf{z}_1, \dots, \mathbf{z}_t) \quad (3)$$

about the pose \mathbf{p}_t of the robot at time t and the position of the landmarks $\mathbf{l}_1, \dots, \mathbf{l}_M$ given all previous motions $\mathbf{u}_1, \dots, \mathbf{u}_t$ and observations $\mathbf{z}_1, \dots, \mathbf{z}_t$. Various approaches to estimate this posterior have been presented in the literature.

In this paper, we address the SLAM problem using the UKF by representing the joint state $(\mathbf{p}_t^T, \mathbf{l}_1^T, \dots, \mathbf{l}_M^T)^T$ with $\langle \mu, \Sigma \rangle$. This is a standard approach which has shown to operate successfully in the past. For convenience, we abbreviate the mean of the robot pose $(\mu_1, \mu_2, \mu_3)^T$ as $(x, y, \theta)^T$. The mean of the j th landmark location $(\mu_{2j+2}, \mu_{2j+3})^T$ is denoted by $(l_x^{[j]}, l_y^{[j]})^T$. Furthermore, we interpret the state transition function h as the robots motion model. In addition, we assume that range and bearing observations $(\rho, \phi)^T$ are given so that we can define a corresponding observation model g . In our work, we initialize new landmarks in a single step following the approach of Bailey [2].

Note that our approach is not limited to UKF or Kalman filter-based approaches and any other method can be applied for addressing the SLAM problem.

C. Monte-Carlo Reinforcement Learning

The basic idea of reinforcement learning [18] is to learn by the interaction with the environment. We consider a dynamical system consisting of an agent and its environment at discrete time steps τ . At each point in time τ , the world is in state $s_\tau \in \mathcal{S}$ and the agent chooses an action $a \in \mathcal{A}$. Then, the world transform into a new state $s_{\tau+1}$ and the agent receives an reward $r_{\tau+1} \in \mathcal{R}$. The goal is to maximize the return

$$R_\tau = \sum_{k=\tau+1}^{\mathcal{T}} r_k, \quad (4)$$

whereas \mathcal{T} is the total number of time steps of one learning episode. The agent is following a policy

$$\pi(s, a) := p(a|s) \quad \forall s \in \mathcal{S} \quad (5)$$

which represents the probability of choosing action a under the assumption of being in state s . Each policy π has a corresponding Q-function

$$Q^\pi(s, a) := E_\pi \{ R_\tau | s_\tau = s, a_\tau = a \} \quad (6)$$

which specifies the expected return R from choosing action a in state s . During the learning process, we would like to approximate the optimal policy $\pi^*(s, a)$ that maximizes the expected return. Therefore, we have to approximate the corresponding Q-function simultaneously.

One way of solving the reinforcement learning problem is based on Monte Carlo methods [3, 17]. Here, we estimate the Q-function as the average return over sample episodes. Initially, the Q-function is initialized with an arbitrary prior. During the training, a soft policy should be used. Thus, it should hold that $\pi(s, a) > 0$ for all possible state-action pairs in order to assure that each state is reachable during the training process. One common soft policy is ϵ -greedy which selects with the high probability of $1 - \epsilon$ the action

$$a^* = \arg \max Q(s, a) \quad (7)$$

that maximizes the expected return and a random action otherwise.

Note that the time index t used in the SLAM setting is not necessarily identical with the discrete time at which the

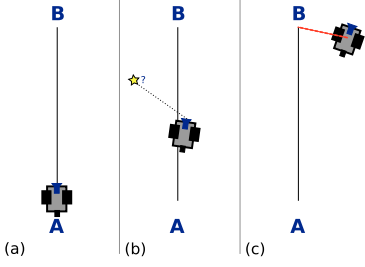


Fig. 1. Illustration of the single-goal navigation task.

reinforcement learning framework has to make decisions. Therefore, we introduced a second time index τ to distinguish both in a sound way.

IV. NAVIGATION TASKS

A. Single-goal Task

Let us consider the following most basic navigation task (see Fig. 1). The robot is located at position A . It is supposed to drive from there to the goal position B . In this example, the robot's motion is affected by a drift. In addition, N landmarks are distributed randomly over the environment. When the robot perceives a new landmark, it has to decide whether it should integrate this landmark in the UKF or not. The UKF has a landmark capacity of M landmarks with $M \ll N$. The goal is to choose the landmarks in a way such that the distance of the final position of the robot $(x_T, y_T)_{\text{true}}^T$ and the target position B is minimized. Hence, we define the reward

$$r_\tau = \begin{cases} -|B - (x_T, y_T)_{\text{true}}^T| & \text{if } \tau = T \\ 0 & \text{else,} \end{cases} \quad (8)$$

as the negative Euclidean distance of the robot's true position to the goal B if the training episode reaches the terminal state s_T ; intermediate rewards $r_1, \dots, r_{\tau-1}$ are set to zero.

B. Round-trip Task

In the round-trip task, the robot is supposed to reach several subgoals. First, it starts at A and it is supposed to drive to B , then back to A . Afterwards, it should target B again and, finally, it should return to A . In this task, a new subgoal is selected as soon as the position estimate of the robot $(x_t, y_t)^T$ is close to it – independent of the robot's true position $(x_t, y_t)_{\text{true}}^T$. In this task, the error in the pose estimate should be minimized over the whole trajectory. For convenience, we specify the return directly as the negative average localization error over the remaining trajectory,

$$R_\tau = -\frac{1}{|T - t(\tau)|} \sum_{t'=t(\tau)}^T \left| \begin{pmatrix} x_{t'} \\ y_{t'} \end{pmatrix}_{\text{true}} - \begin{pmatrix} x_{t'} \\ y_{t'} \end{pmatrix} \right|, \quad (9)$$

whereas $t(\tau)$ specifies the time when the τ th decision is made and T is the time when the robot reaches its final destination. To simplify things for the second task, landmark selection is only allowed while the robot moves from A to B the first time. The round-trip task is more complex than the previous one. However, it is worth considering since it focuses on the loop-closing problem of SLAM and has a higher a practical relevance than the single-goal task.

V. NAVIGATION AND LANDMARK SELECTION

A. Motion Control

The robot is steered towards the subgoals using a straightforward controller. An appropriate translational acceleration $\dot{\omega}_t$ and rotational acceleration \dot{v}_t is selected based on the current estimate of the robot pose \mathbf{p}_t , the translational velocity ω_t and rotational velocity v_t .

B. Learning Landmark Selection Policies

In order to learn landmark selection policies with Monte Carlo reinforcement learning, we need to define the state space \mathcal{S} and the action space \mathcal{A} . In addition to that, we need to find an appropriate representation for the continuous Q -function.

1) *State Space*: The available state information consists of the UKF state $\langle \mu, \Sigma \rangle$ and the current range and bearing observation $(\rho, \phi)^T$. This full information would lead to an high-dimensional state space so that a successful learning is impractical. It is therefore desirable to reduce the space while preserving as much as possible of the relevant information. This can be achieved by defining features that summarize the essential information. One of the features is the position of the potentially new landmark,

$$\begin{pmatrix} l_x^{[\text{new}]} \\ l_y^{[\text{new}]} \end{pmatrix} = \begin{pmatrix} x_t + \rho \cos(\phi + \theta_t) \\ y_t + \rho \sin(\phi + \theta_t) \end{pmatrix}, \quad (10)$$

according to the current range and bearing observation $(\rho, \phi)^T$ and the robot's pose estimate $(x_t, y_t, \theta_t)^T$. Additionally, we define the following five features:

1. Estimated distance to subgoal B,

$$d_{\text{est}} = |B - (x_t, y_t)^T|, \quad (11)$$

2. Number of landmarks integrated in the UKF,

$$m = |\{j \in M : \Sigma_{2j+2} < \infty \wedge \Sigma_{2j+3} < \infty\}|, \quad (12)$$

where Σ_{2j+2} and Σ_{2j+3} are the variances of the j th landmark in the x and y direction.

3. Yaw angle to potential new landmarks ϕ ,
4. Distance of the potentially new landmark to the closest landmark already integrated,

$$d_l = \min_{\substack{j \in L \\ \text{with } \Sigma_{2j+2} < \infty \\ \wedge \Sigma_{2j+3} < \infty}} \left| \begin{pmatrix} l_x^{[j]} \\ l_y^{[j]} \end{pmatrix} - \begin{pmatrix} l_x^{[\text{new}]} \\ l_y^{[\text{new}]} \end{pmatrix} \right|, \quad (13)$$

5. Uncertainty of the robot pose $\Sigma^{3 \times 3}$ in terms of its entropy,

$$H = \ln(\sqrt{(2\pi e)^3 |\Sigma^{3 \times 3}|}). \quad (14)$$

The first of these features summarizes the robot position (x_t, y_t) . The landmark positions $\mathbf{l}_1, \dots, \mathbf{l}_M$ are summarized by the fourth feature while the new observation $(\rho, \phi)^T$ is represented by the third feature as well as fourth one. The covariance Σ_t is comprised by the second feature and the fifth one.

In the following, we will consider three different variants of the learning approaches. The first approach only relies on a two dimensional state space (first and second feature), the second one uses an four dimensional feature space (first to fourth feature) and the third one uses five dimensions (all five features).

2) *Function Approximation*: Since the state space of the features is continuous (with the exception of the second dimension), we need to estimate the Q-function with some function approximator. In our current implementation, we use k -nearest neighbor (k -NN) regression [16]. Training points – i.e. state/action values (s, a) which are each labeled with a return R – are efficiently stored in set of kd-trees [4, 1]. The j th kd-tree represents the returns from choosing an action a_j in a given state s . If a query (s', a') is performed, the k nearest data points to the query point s' (w.r.t. Euclidean distance) are selected from the appropriate kd-tree. The return is estimated as unweighted average over the corresponding R -values. If less then k_{min} data points can be found within a fixed radius around the query point, some prior R_{prior} is returned. In our current implementation, we set $k = 50$ to reflect the high amount of noise in our training data; k_{min} is set to 10. The k -NN regression approach has the advantage over the common grid-based discretization methods that it has a high degree of generalization in areas where the density is low and it is precise in regions where the data points are dense. In contrast to non-linear models such as neural networks [14], no over-fitting occurs. As opposed to other regression techniques, in which the model is also expressed directly in terms of their training data such as Gaussian Processes [13], k -NN regression is very fast. Even with hundred thousands of data points, a query can be performed in a few milliseconds. An efficient evaluation is essential for learning in practice, since the regression has to be carried out frequently. In various tests, we could not reveal a significant benefit from using Gaussian process regression over k -NN regression for reinforcement learning in our domain. Due to space restrictions, these experiments are omitted in the experimental section.

3) *Action Selection*: In our learning problem, the action is a binary decision:

$$\mathcal{A} = \{a_{reject}, a_{accept}\} \quad (15)$$

This means that either the potential new landmark is chosen or not. In order to boost the training, a variant of ϵ -greedy is used:

$$\pi(s) = \begin{cases} \arg \max Q(s, a) & \text{if } Q(s, a_{reject}) \neq Q(s, a_{accept}) \\ & \text{and } \chi_1 < 1 - \epsilon \\ a_{accept} & \text{if } [Q(s, a_{reject}) = Q(s, a_{accept}) \\ & \text{or } \chi_1 \leq \epsilon] \text{ and } \chi_2 < \frac{M}{N_{visible}} \\ a_{reject} & \text{else} \end{cases} \quad (16)$$

Here, χ_1 and χ_2 are uniform random samples between zero and one; $N_{visible}$ is the expected number of visible landmarks in one training episode. Thus, in the beginning of the training – when $Q(s, a_{reject}) = Q(s, a_{accept}) = R_{prior}$

in most cases – it is ensured that landmarks are selected over the whole trajectory. Neither landmarks in the beginning of the episode nor landmarks in the end are preferred. If standard ϵ -greedy is used, a_{accept} and a_{reject} would be chosen with a probability of 0.5 each. Thus, depending on the values for the landmark capacity M and expected number of visible landmarks $N_{visible}$ it could happen that either all landmarks are selected in the beginning of the episode or that considerably fewer landmarks than M are selected. Either would lead to a slow convergence rate.

To sum up, we use a learning approach for landmark selection based on Monte-Carlo reinforcement learning and k -NN regression. The state space is compactly represented by five features and the action is a binary decision.

C. Generalization

Until now, we considered an approach to learn a selection policy in unknown environments, but for a specific scenario. However, it is desirable to train a policy in one scenario and then apply this policy in another setting. Important parameters of a training scenario are the number N of landmarks in the environment and the landmark capacity M of the UKF. To generalize, it is important to have a scenario-independent state space representation. For instance, instead of the number of landmark integrated in the UKF m , we need to speak about the percentage of landmarks $\frac{m}{M}$.

D. Deletion of Landmarks

In the Kalman filter framework, it is possible to delete already integrated landmarks. This can be done without affecting the statistical consistency by removing the appropriate value from the mean vector and the corresponding rows and columns from the covariance matrix [6]. If we want to allow deletion, we must extend our action set \mathcal{A} . Since deletion is only useful if we replace the deleted landmark with a new one, we propose the following action set

$$\mathcal{A} = \{a_{reject}, a_{replace_1}, \dots, a_{replace_M}\}. \quad (17)$$

The deletion of landmarks might be particular interesting in connection with the round-trip task where those landmarks should not be replaced which facilitate loop closure.

VI. EXPERIMENTS

A. Single-goal Task in Simulation

We evaluate the performance of our learning procedure for the single-goal task in a simulated environment. We choose an environment where N landmarks are randomly distributed in a $30m \times 60m$ area. The distance between the start position A and the goal B is set to $44m$. We train our policy for 2,000 episodes. In each episode, landmarks are randomly re-distributed. We compare the trained policies with two heuristics. The first one is the *M-first heuristic* which simply integrates the M first landmarks that are observed. An apparently better policy is the *equidistant heuristic*. With this heuristic, the robot only integrates a new landmark after it has driven a certain distance so that the

Test scenario M / N	Policy trained in scenario						equidistant heuristic
	5 / 100	5 / 50	10 / 100	10 / 50	15 / 100	15 / 50	
5 / 100	11.5 ± 0.64	12.2 ± 0.68	14.9 ± 1.19	15.6 ± 1.36	16.6 ± 1.05	17.4 ± 0.99	13.2
5 / 50	10.7 ± 0.42	10.5 ± 0.49	11.6 ± 0.50	12.2 ± 0.52	12.7 ± 0.49	13.5 ± 0.62	13.6
10 / 100	6.9 ± 0.54	6.3 ± 0.47	6.8 ± 0.39	7.6 ± 0.61	8.4 ± 0.6	9.7 ± 0.97	8.5
10 / 50	8.2 ± 1.05	7.2 ± 0.48	7.1 ± 0.45	6.8 ± 0.26	7.1 ± 0.47	7.2 ± 0.35	9.6
15 / 100	5.9 ± 0.87	5.1 ± 0.36	4.9 ± 0.35	5.1 ± 0.24	5.1 ± 0.22	6.0 ± 0.46	6.6
15 / 50	8.0 ± 1.24	6.9 ± 0.71	6.5 ± 0.76	5.9 ± 0.41	6.0 ± 0.22	5.6 ± 0.28	7.5

TABLE I

HIGH DEGREE OF GENERALIZATION IN THE SINGLE-GOAL TASK. THE MEAN ERROR OVER TEN TRAINING RUNS AND THE CORRESPONDING STANDARD DERIVATION IS SHOWN. ALL POLICIES MARKED BOLD ARE SIGNIFICANTLY BETTER THAN THE EQUIDISTANT HEURISTIC ($\alpha = 0.05$).

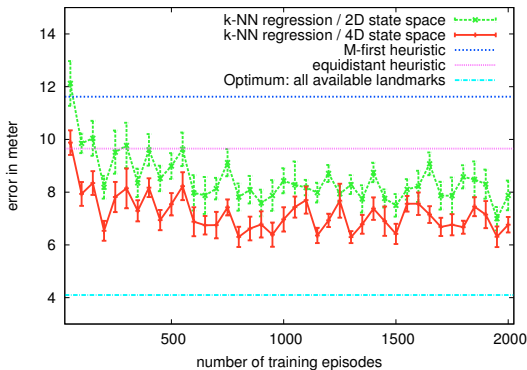


Fig. 2. Evolution of the error during training using k -NN regression. The mean errors over 10 training runs as well as the corresponding 95%-confidence intervals are plotted in the graph. In addition to that, the average performance of two different heuristics over 2000 episodes is shown. Furthermore, the figure shows the average performance of navigation under optimal conditions where all visible landmark are integrated in the UKF.




Fig. 3. Average performance of the trained policies and heuristics w.r.t. 1000 test episodes. For the trained policies, the mean over the ten training runs as well as the corresponding 95%-confidence interval is shown.

landmarks are approximately uniformly distributed over the whole trajectory (similar to [6]).

For the learning, we use k -NN regression with a two-, a four-, and a five-dimensional state space (see Section V-B). For comparison, a learning approach using a grid-based discretization is performed. At first, we consider an UKF with a landmark capacity of $M = 10$ and an environment with $N = 50$ landmarks. Fig. 2 shows the evolution of the error of different learning approaches using k -NN regression during the training. The error is defined by the Euclidean distance of the robot to the goal B . For each learning approach, ten training runs are performed. Each trained policy and heuristic is evaluated in 1,000 different environments (see Fig. 3). The one-sample t-test with a significance level of $\alpha = 0.05$ shows that all four learning approaches are significantly better than the equidistant heuristic. Furthermore, it is shown using a two-sample t-test that k -NN regression with a four




Fig. 4. Pioneer 2-DX8 robot with upward-looking camera and SICK laser range scanner (left) and detected visual landmarks on the ceiling (right).

dimensional state space leads to a significant smaller error than grid-based discretization with four dimensions as well as k -NN with two dimensions. Thus, the third feature, which is the distance d_l of a new landmark to landmarks already integrated, and the fourth one, which is the angle ϕ to the new landmark, seem to include relevant information which are not encoded in the first two dimension of the state space. Further experiments revealed that indeed both features are essential. However, we were not able to show that there is any benefit from including the fifth feature, the entropy H of the robot's pose. Even with a significance level of $\alpha = 0.25$, the t-test did not reveal a difference between the learning approach using the four dimensional state space and the one using five dimensions. A qualitative comparison in an example environment between the learned policy and the heuristics is shown in the first half of the accompanying video submission.

In order to evaluate how good the trained policies generalize, we trained and tested a policy in environments with $N = 50$ as well as $N = 100$ landmarks. In addition, we use UKFs with a capacity M of five, ten, and 15 landmarks. Tab. I illustrates the high degree of generalization of our learning approach. For instance, if we perform a training in a setting with $N = 50$ and $M = 5$, we see that the trained policy leads to significantly better results than the equidistant heuristic in all six test scenarios. This indicates that our approach generalized over different landmark densities which is similar to environments of different scale and sensor range.

B. Single-goal Task Performed in a Real World Experiment

Furthermore, we evaluated our learning approach in a laboratory environment. Visual markers [8] have been randomly attached to the ceiling in a $2.5m \times 5m$ area. Our used robot, a Pioneer 2DX-8, is equipped with an upward-looking camera and a SICK laser range scanner (see Fig. 4). The camera is used for the experiments observing landmarks

at the ceiling whereas the laser is used for (near) ground truth evaluation. Since the odometry of the robot was too accurate in the limited space in which we carried out the experiment, we added a rotational bias of 0.1 rad per meter. It is impractical to train the policy in the real-world because this would not only require to perform hundreds of training episodes but also to install different landmark distributions for each training episode. Thus, we trained the policy in simulation and tested it in the real-world setting. We also compared the trained policy to the equidistant heuristic. Both, the trained policy as well as the equidistant heuristic were tested ten times. The trained policy results in an error of 0.50 ± 0.08 whereas the equidistant heuristic leads to an error of 0.66 ± 0.07 . Hence, the trained policy is significantly better than the equidistant heuristic (w.r.t. a t-test with $\alpha = 0.05$).

C. Round-trip Task

The performance of our learning procedure for the round-trip task is evaluated in a simulated environment with $N = 50$ landmarks. It was trained using k -NN regression with a four dimensional state space over ten training runs. Here, the error is defined as the average localization error over the whole trajectory. Again, we compare our learning with the equidistant heuristic. Tab. II shows that the learned policy is significantly better than the heuristic. Furthermore, it is shown that we were able to generalize over the UKF capacity M . In the second half of the accompanying video, a qualitative evaluation is given.

D. Round-trip Task with Landmark Deletion

Finally, we analyzed a variant of the round-trip task that, compared to the previous experiments, also allows the deletion of landmarks. It should be noted that this scenario is significantly more complex compared to the previous tasks. For instance, the state space must be extended in order to represent the already integrated landmarks. In initial experiments, we figured out that good strategies keep a set of landmarks fixed for re-localization and perform an incremental pose correction with set of frequently replaced landmarks. Further investigations are currently ongoing.

VII. CONCLUSION

In this paper, we presented an novel approach for landmark selection in unknown environments using reinforcement learning. The ability of a mobile robot to incorporate a landmark into its belief or to discard it allows for efficient robot navigation under computational constraints. The presented method is able to determine which landmark is valuable for the robot to efficiently solve its current navigation task. This is especially important for robots with restricted resources. We demonstrated by a series of real world and simulation experiments that the learned policies outperform handcrafted heuristics. Furthermore, we showed that a learned policy has a high degree of generalization since it can be applied in different environments with changed underlying parameters.

Despite these encouraging results, there is space for further optimizations. One interesting aspect is the possibility to

Test M	Policy trained in scenario			equidistant heuristic
	5	10	15	
5	3.28 ± 0.15	4.44 ± 0.41	5.06 ± 0.37	3.86
10	2.38 ± 0.09	2.37 ± 0.09	2.43 ± 0.06	2.85
15	2.30 ± 0.13	2.24 ± 0.17	2.23 ± 0.06	2.55

TABLE II

ROUND-TRIP TASK. ALL POLICIES MARKED BOLD ARE SIGNIFICANTLY BETTER THAN THE EQUIDISTANT HEURISTIC ($\alpha = 0.05$).

delete an already incorporated landmark. First experiments indicate that improvements can be made although the problem is substantially more complex.

REFERENCES

- [1] S. Arya and D.M. Mount. Algorithms for fast vector quantization. In *Proc. of the IEEE Data Compression Conference (DDC'93)*, pages 381–390, 1993.
- [2] T. Bailey. *Mobile Robot Localisation and Mapping in Extensive Outdoor Environments*. PhD thesis, University of Sydney, 2002. pages 22–23.
- [3] A. Barto and M. Duff. Monte Carlo matrix inversion and reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 687–694, 1994.
- [4] J.L. Bentley. K-d trees for semidynamic point sets. In *Proc. of the 6th ACM Symposium on Computational Geometry*, pages 187–197, 1990.
- [5] Kwok C. and Fox D. Reinforcement learning for sensing strategies. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'04)*, Sendai, Japan, 2004.
- [6] G. Dissanayake, H. Durrant-Whyte, and T. Bailey. A computationally efficient solution to the simultaneous localisation and map building (SLAM) problem. In *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA'00)*, pages 1009–1014, 2000.
- [7] S.J. Julier and J.K. Uhlmann. A new extension of the Kalman filter to nonlinear systems. In *Proceedings of the Int. Symp. on Aerospace/Defense Sensing, Simulation and Controls*, 1997.
- [8] H. Kato and M. Billinghurst. Marker tracking and HMD calibration for a video-based augmented reality conferencing system. In *Proc. of the Int. Workshop on Augmented Reality (IWAR'99)*, 1999.
- [9] R. Lerner, E. Rivlin, and I. Shimshoni. Landmark selection for task-oriented navigation. *IEEE Transaction on Robotics*, 23(3), 2007.
- [10] Bryson M. and Sukkarieh S. Active airborne localisation and exploration in unknown environments using inertial SLAM. In *Proc. of the IEEE Aerospace Conference*, 2006.
- [11] P.S. Maybeck. *The Kalman filter: An introduction to concepts*. In *Autonomous Robot Vehicle*. Springer Press, 1990.
- [12] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proc. of the National Conf. on Artificial Intelligence (AAAI'02)*, pages 593 – 598, 2002.
- [13] C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, USA, 2006.
- [14] R. Rojas. *Neural Networks: A Systematic Introduction*. Springer Press, 1996.
- [15] P. Sala, R. Sim, A. Shokoufandeh, and S. Dickinson. Landmark selection for vision-based navigation. *IEEE Transaction on Robotics*, 22(2), 2006.
- [16] G. Shakhnarovich, T. Darrell, and P. Indyk. *Nearest-Neighbor Methods in Learning and Vision: Theory and Practice*. MIT Press, Cambridge, MA, USA, 2006.
- [17] S.P. Singh, R.S. Sutton, and P. Kaelbling. Reinforcement learning with replacing eligibility traces. In *Machine Learning*, volume 22, pages 123–158, 1996.
- [18] R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*. MIT Press, Cambridge, MA, USA, 1998.
- [19] S. Thrun, Y. Liu, D. Koller, A.Y. Ng, Z. Ghahramani, and H. Durrant-Whyte. Simultaneous localization and mapping with sparse extended information filters. *Int. Journal of Robotics Research*, 23, 2004.
- [20] S. Zhang, L. Xie, and M.D. 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'05)*, 2005.