

# Near-optimal Landmark Selection for Mobile Robot Navigation

Maximilian Beinhofer    Jörg Müller    Wolfram Burgard

**Abstract**—The ability to accurately localize themselves is a fundamental pre-condition for service robots designed to carry out navigation and transportation tasks. Because of the high degree of dynamics in populated and real-world environments, often artificial landmarks are used to achieve the desired accuracy in localization. In this paper we consider the problem of optimally placing landmarks for robots navigating frequently on similar trajectories. Our method maximizes conditional mutual information of the states of the robot given the landmark observations. It uses a greedy algorithm which approximates the solution of the NP-hard maximization problem. For this algorithm, we derive a tight constant-factor bound on the approximation error. We furthermore evaluate the selected landmark sets in extensive experiments carried out both in simulation and with a real robot.

## I. INTRODUCTION

One of the fundamental problems in mobile robot navigation is localization. Especially in industrial applications, where a high degree of robustness is required even when the environments are dynamic, a popular solution to this problem is the use of active beacons [4], [10] or artificial landmarks [5], [17].

However, the number of landmarks to be used is usually restricted, e.g., by the limited computational capabilities of embedded systems or the costs of active beacons. Therefore, one is faced with the problem of how to place the landmarks or beacons efficiently so that the uncertainty of the robot in its position estimate is minimized. Especially in industrial environments robots often have to repeatedly execute a fixed number of trajectories. The locations of the landmarks should be chosen depending on these trajectories. In practice, the landmark positions are often manually selected [4].

In this paper, we consider landmark selection as an optimization problem. For one given trajectory (like the one in Fig. 1), we place a given number  $n$  of landmarks so that the mutual information [3] between all states of the robot along the trajectory and all landmark observations is maximized. The goal is to arrive at a landmark configuration that yields the highest information gain. As shown by Krause and Guestrin [7], this problem is NP-hard. We introduce a greedy algorithm which approximates the solution of the problem. Because of the non-linear system dynamics, this algorithm employs Monte Carlo simulation to evaluate the objective function. Using the concept of submodularity [9], we prove that our algorithm finds near-optimal sets of landmarks. We evaluate the selected landmark sets in simulation

This work has partly been supported by the German Research Foundation (DFG) within the Research Training Group 1103. All authors are with the Department of Computer Science at the University of Freiburg, Germany. beinhofe@informatik.uni-freiburg.de

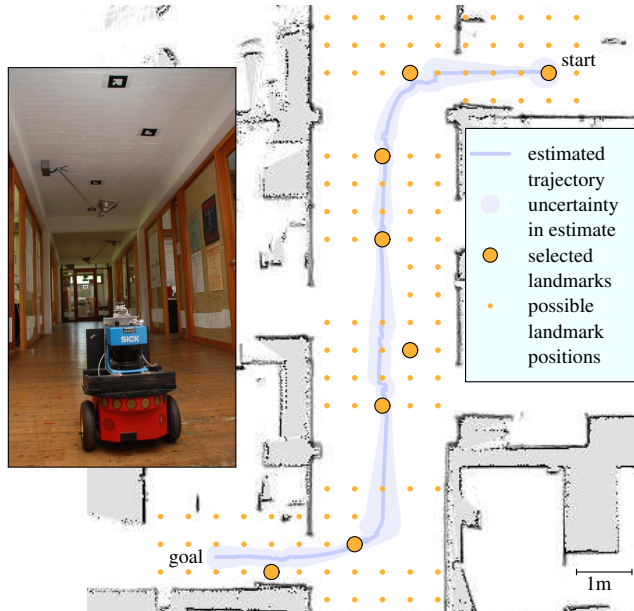


Fig. 1. Eight landmarks selected by our algorithm. One execution of the trajectory for which the landmarks were selected is plotted in blue together with its  $2\sigma$  uncertainty. The picture shows the Pioneer P3-DX robot used in the experiments. It is equipped with a SICK LMS laser range finder (only used for reference) and a webcam pointing upwards. The markers on the ceiling correspond to the landmarks shown in the graph.

and on a real robot, executing the trajectory displayed in Fig. 1. The results of the experiments show that our algorithm significantly outperforms other approaches.

This paper is organized as follows. After discussing related work in the following section, we introduce the landmark selection problem in Section III. We then present our approximation algorithm and show a tight bound on the approximation error in Section IV. Afterwards, in Section V, we discuss different models for controlling the robot and describe their influence on the properties of the approximation. Finally we present extensive experiments in which we evaluate our algorithm in simulation as well as with a real robot.

## II. RELATED WORK

In the past, a substantial amount of work has been carried out in the context of landmark selection. The majority of approaches can be subdivided into two categories, namely online and offline selection. In the online selection task, the robot at each time step has to select the subset of visible features that yields the most information. Usually no artificial landmarks are introduced in this setting. The second class is that of offline selection, which we also consider in this paper. The task is to select the most useful landmarks before the robot starts operating. Note that the offline selection task

corresponds to the sensor placement task, where sensors are placed in the environment to observe the (blind) robot to estimate its position.

One example of an online method is Thrun’s approach [13]. In the context of a probabilistic localization framework, it considers the average posterior localization error, a value that depends on the actual belief of the robot, and uses neural networks to extract landmarks from sensor snapshots. Zhang *et al.* [18] and Strasdat *et al.* [12] considered the landmark selection task in the context of the simultaneous localization and mapping (SLAM) problem. Whereas Zhang *et al.* [18] select at every time step observations which minimize the entropy of the resulting posterior distribution, Strasdat *et al.* use reinforcement learning to determine a landmark selection policy.

An approach that belongs to the class of offline selection methods is that of Sala *et al.* [11] which considers the problem independently of the trajectory and in terms of the art gallery problem. Their algorithm selects landmarks so that from every position in the environment, the robot has a clear line of sight to at least  $k \geq 1$  landmarks. Likewise, Meyer-Delius *et al.* [8] present an approach that is independent of the trajectory taken by the robot. They increase the localization accuracy of a system already equipped with a landmark-independent sensor (e.g., a laser range finder) by placing landmarks so as to maximize the average uniqueness of measurements at all poses throughout the map (see Section VI-A). Jourdan and Roy [6] propose an algorithm for placing sensors on the boundaries of buildings, which minimizes the position error bound (PEB), a lower bound on localization accuracy. The work of Vitus and Tomlin [16] is probably closest to our approach as it performs trajectory-dependent sensor (landmark) placement considering the belief of the robot about its state. They minimize the sum of the traces of the covariance matrices of all states.

In contrast to previous methods we also derive a tight bound on the approximation error of our greedy algorithm following the concept laid out by Krause and Guestrin [7]. Besides other results, they prove submodularity of information gain for special sets of discrete random variables in graphical models. With this property, they can give a  $(1 - 1/e)$  approximation guarantee for the greedy algorithm to find the set of random variables which maximizes information gain. In this paper, we expand their concept to the case of continuous random variables and differential mutual information.

### III. LANDMARK SELECTION AS AN OPTIMIZATION PROBLEM

As described in the previous section, we consider landmark selection as a trajectory-dependent problem. The evolution of the state  $X_t$  at time  $t \in [1, T]$  of the robot executing a given trajectory can be described by a Hidden Markov Model (HMM). For the task of placing landmarks, the HMM has to explicitly include all possible landmark positions  $\mathcal{V} = L^{1:N}$  and the separate observations  $Z_t^{1:N}$  at time  $t$  of the different landmarks. Without the dashed arrows, Fig. 2 shows the




Fig. 2. The dynamic Bayesian network for the localization of a mobile robot. It characterizes the evolution of controls  $U$ , states  $X$ , measurements  $Z$  depending on the deterministic landmarks  $L$ . The dashed arrows model the additional dependency for an external controller.

dynamic Bayesian network that describes the HMM. In this network,  $U_t$  is the control command at time  $t$ , which for simplicity is typically considered to be randomly chosen in the localization task [14].

In order to explicitly calculate probabilities in this HMM, we need to specify several properties of the robot and the trajectory. More precisely, we take the motion model and sensor model of the robot including the maximum range of the sensor and the control policy and the desired trajectory as inputs to the optimization procedure. We also consider the number  $n$  of landmarks to be placed as given (e.g., by the constrained memory of the robot), but one could also use a threshold as in [8] to determine  $n$ .

For every environment and every kind of landmark detecting sensor, there exists a subspace of the environment in which it is possible to place landmarks. For our landmark selection approach, we use a discrete representation  $\mathcal{V} = \{L^1, \dots, L^N\}$  of this subspace. On the power set of this discrete set, we define an objective function  $F : \mathcal{P}(\mathcal{V}) \rightarrow \mathbb{R}$ . For every subset  $\mathcal{A} \in \mathcal{V}$ ,  $F(\mathcal{A})$  describes the value of information gained by placing landmarks in all points in  $\mathcal{A}$ . Given a constrained number  $n$  of landmarks the robot can use for localization, our goal is to find

$$\mathcal{A}^* = \operatorname{argmax}_{\mathcal{A} \subseteq \mathcal{V}; |\mathcal{A}| \leq n} F(\mathcal{A}). \quad (1)$$

As objective function  $F$ , we consider conditional mutual information (8.5 in [3]) in the Hidden Markov Model. In the HMM, placing landmark  $L^k$  is equivalent to observing the random variables  $Z_{1:T}^k$ . Note that  $Z_t^k$  can also take one special value indicating that  $L^k$  is not visible at time  $t$ . Using the notation  $Z_{1:T}^{\mathcal{A}} := \{Z_{1:T}^i \mid L^i \in \mathcal{A}\}$ , the objective function has the following form:

$$\begin{aligned} F(\mathcal{A}) &= I(X_{1:T}; Z_{1:T}^{\mathcal{A}} \mid U_{1:T}, L^{1:N}) \\ &= h(X \mid U, L^{1:N}) - h(X \mid Z^{\mathcal{A}}, U, L^{1:N}) \\ &= h(X \mid U) - h(X \mid Z^{\mathcal{A}}, U). \end{aligned} \quad (2)$$

Note that the positions  $L^k$  of the landmarks are deterministic and globally known for all landmarks in  $\mathcal{V}$ , so conditioning on them does not change the considered probability distributions. Also, in (2) and in the following, we omit

the indices  $1 : T$  for convenience. Because in general  $X$ ,  $U$  and  $Z$  are sets of random variables which live on continuous probability spaces,  $h$  is the differential entropy. The objective function in (2) describes how the entropy of the joint probability distribution of all the states of the robot on the trajectory is reduced by the selected set of landmarks (resp. their observations). As the entropy can be viewed as a measure of uncertainty for the underlying distribution, this objective function is an intuitive choice to us.

#### IV. APPROXIMATION ALGORITHM

As a result of Theorem 4 in this section, the problem defined in (1) is **NP**-hard. Since there exists no deterministic polynomial time algorithm to determine the exact maximum  $F(\mathcal{A}^*)$ , we have to be content with an approximation algorithm. Selecting the landmark subset which maximizes conditional mutual information is equivalent to finding

$$\mathcal{A}^* = \operatorname{argmin}_{\mathcal{A} \subseteq \mathcal{V}; |\mathcal{A}| \leq n} h(X | Z^{\mathcal{A}}, U). \quad (3)$$

This follows directly from the definition of the objective function in (2). Algorithm 1 approximates  $\mathcal{A}^*$  greedily using this equivalence. The complexity of the algorithm is in

---

#### Algorithm 1 Greedy approximation algorithm

---

**Input:**  $\mathcal{V}, n$

**Output:**  $\mathcal{A}$

$\mathcal{A} := \emptyset$

**for**  $i := 1$  to  $n$  **do**

**for all**  $L \in \mathcal{V}$  **do**

$h := h(X | Z^{\mathcal{A} \cup \{L\}}, U)$

**end for**

$L^* := \operatorname{argmin}\{h | L \in \mathcal{V}\}$

$\mathcal{A} = \mathcal{A} \cup \{L^*\}$

**end for**

**return**  $\mathcal{A}$

---

$\mathcal{O}(n \cdot |\mathcal{V}|)$ , which makes it applicable even for large values of  $|\mathcal{V}|$ . Note that this algorithm could be easily extended to the case of maximizing mutual information over several possibly weighted trajectories.

##### A. Submodularity of conditional mutual information

To get a constant factor approximation guarantee for Algorithm 1, we use the same concept as Krause and Guestrin in [7], namely that of submodularity.

Nemhauser *et al.* [9] show a factor  $(1 - 1/e)$  approximation guarantee for greedy algorithms for objective functions which satisfy the following three properties:

- $F(\emptyset) = 0$
- submodularity, i.e.,  $\forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}, L \notin \mathcal{B}$  :  
 $F(\mathcal{A} \cup \{L\}) - F(\mathcal{A}) \geq F(\mathcal{B} \cup \{L\}) - F(\mathcal{B})$
- monotonicity, i.e.,  $\forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$  :  $F(\mathcal{A}) \leq F(\mathcal{B})$ .

We chose conditional mutual information as objective function specifically to have these properties. In the landmark selection application, submodularity is an intuitive property. It means that the increase in information about the state of

the robot resulting from adding a new landmark is lower when more landmarks have already been placed.

All proofs in this section follow the concepts laid out by Krause and Guestrin [7]. As can be seen in (2),  $F(\emptyset) = 0$  holds obviously.

*Theorem 1:*  $F(\mathcal{A}) = I(X; Z^{\mathcal{A}} | U)$  is submodular.

*Proof:* In the Hidden Markov Model described above, it holds that for every  $\mathcal{A}$ ,  $Z^{\mathcal{A}}$  is a set of random variables which are independent conditioned on  $X$ . This Markov property follows directly from the rules of d-separation [2] (see Fig. 2). Using this and the definition of conditional entropy [3], we get for every  $\mathcal{A}$ :

$$\begin{aligned} I(X; Z^{\mathcal{A}} | U) &= h(X | U) - h(X | Z^{\mathcal{A}}, U) \\ &= h(X | U) - h(Z^{\mathcal{A}} | X, U) - h(X, U) + h(Z^{\mathcal{A}}, U) \\ &= h(X | U) - h(X, U) + h(Z^{\mathcal{A}}, U) \\ &\quad - \sum_{L' \in \mathcal{A}} h(Z^{\{L'\}} | X, U). \end{aligned} \quad (4)$$

In (4),  $h(X | U) - h(X, U)$  does not depend on  $\mathcal{A}$  and  $\sum_{L' \in \mathcal{A}} h(Z^{\{L'\}} | X, U)$  adds the same value to both sides of the submodularity inequation. So to prove submodularity of  $I(X; Z^{\mathcal{A}} | U)$ , it suffices to prove submodularity of  $h(Z^{\mathcal{A}}, U)$ . We do this by considering the definition of submodularity.  $\forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}, L \notin \mathcal{B}$  :

$$\begin{aligned} h(Z^{\mathcal{A} \cup \{L\}}, U) - h(Z^{\mathcal{A}}, U) &= \\ h(Z^{\{L\}} | Z^{\mathcal{A}}, U) &\geq h(Z^{\{L\}} | Z^{\mathcal{B}}, U) \\ &= h(Z^{\mathcal{B} \cup \{L\}}, U) - h(Z^{\mathcal{B}}, U). \end{aligned}$$

The inequation holds because of the ‘‘conditioning reduces entropy’’-principle [3], which states that  $h(A | B) \leq h(A)$  for all  $A, B$ . ■

*Theorem 2:*  $F(\mathcal{A}) = I(X; Z^{\mathcal{A}} | U)$  is monotonically increasing in  $\mathcal{A}$ .

*Proof:* Using the ‘‘conditioning reduces entropy’’-principle [3] and the conditional independence property in the HMM, it holds that  $\forall \mathcal{A} \subseteq \mathcal{V}, L \notin \mathcal{A}$  :

$$\begin{aligned} F(\mathcal{A} \cup \{L\}) - F(\mathcal{A}) &\stackrel{(2)}{=} h(X | Z^{\mathcal{A}}, U) - h(X | Z^{\mathcal{A} \cup \{L\}}, U) \\ &\stackrel{(4)}{=} h(Z^{\mathcal{A} \cup \{L\}}, U) - h(Z^{\mathcal{A}}, U) - h(Z^{\{L\}} | X, U) \\ &= h(Z^{\{L\}} | Z^{\mathcal{A}}, U) - h(Z^{\{L\}} | X, U) \quad (5) \\ &\geq h(Z^{\{L\}} | X, Z^{\mathcal{A}}, U) - h(Z^{\{L\}} | X, U) \\ &= h(Z^{\{L\}} | X) - h(Z^{\{L\}} | X) = 0. \quad (6) \end{aligned}$$

Theorems 1 and 2 provide the properties that are required in Nemhauser *et al.* [9] for the following bound:

*Theorem 3:* For the result  $\mathcal{A}_{\text{greedy}}$  of Algorithm 1, the following approximation guarantee holds:

$$F(\mathcal{A}_{\text{greedy}}) \geq (1 - 1/e) \max_{\mathcal{A} \subseteq \mathcal{V}; |\mathcal{A}| \leq n} F(\mathcal{A}). \quad (7)$$

For the objective function considered in this paper, this Nemhauser bound is a tight bound:

*Theorem 4:* The optimization problem defined in (1) is not approximable in polynomial time within a constant factor better than  $(1 - 1/e)$ , unless  $\mathbf{P} = \mathbf{NP}$ .

*Proof:* The proof of Theorem 9 in [7] also covers this case. ■

Theorem 4 states that no other polynomial time algorithm can have a better worst case approximation guarantee than the greedy algorithm.

### B. Calculation of Entropy of the Full Posterior in the Bayes Filter

Throughout this paper, we consider the problem of estimating the pose of a robot relative to a given set of landmarks  $\mathcal{A}$  using a Bayes filter. The key idea of the Bayes filter approach is to maintain a probability density  $p(x_t | z_{1:t}, u_{1:t})$  (called belief) of the random variable  $X_t$  describing the pose of the robot at time  $t$  given all landmark observations  $z_{1:t}$  and control inputs  $u_{1:t}$  up to time  $t$ . This probability is calculated recursively using the Bayesian filtering scheme

$$p(x_t | z_{1:t}, u_{1:t}) = \eta_t p(z_t | x_t) \cdot \int p(x_t | u_t, x_{t-1}) p(x_{t-1} | z_{1:t-1}, u_{1:t-1}) dx_{t-1}. \quad (8)$$

Here,  $\eta_t$  is a normalizer that ensures that  $p(x_t | z_{1:t}, u_{1:t})$  integrates up to 1 over all  $x_t$ . The term  $p(x_t | u_t, x_{t-1})$  is the motion model and  $p(z_t | x_t)$  the sensor model, respectively.

In the greedy approximation Algorithm 1, we need to evaluate conditional differential entropies of the form  $h(X | Z^{\mathcal{A}}, U)$ . Because these terms contain the highly correlated random variables  $Z_{1:T}^{\mathcal{A}}$  and  $U_{1:T}$ , their evaluation includes solving a  $(T \cdot (\dim(Z_{1:T}^{\mathcal{A}}) + \dim(U_{1:T})))$ -dimensional integral over a term for which in general no antiderivative exists. Therefore, we apply Monte Carlo simulation and obtain an approximation  $\hat{h}$  of the entropy  $h(X | Z^{\mathcal{A}}, U)$  as the average over  $N$  simulated runs of the robot. In each run of the simulation, instead of using the probability density  $p(x_{1:T} | z_{1:T}^{\mathcal{A}}, u_{1:T})$  to calculate the entropy of the run, we use the factorization (see (11.9) in [14])

$$p(x_{1:T} | z_{1:T}^{\mathcal{A}}, u_{1:T}) = \prod_{t=1}^T \eta_t p(z_t^{\mathcal{A}} | x_t) p(x_t | x_{t-1}, u_t). \quad (9)$$

In each single run, this factorization can be efficiently evaluated via recursion.

The Monte Carlo approximation  $\hat{h}$  of  $h$  of course increases the approximation error in Algorithm 1. However, the approximation guarantee (Theorem 3) can be extended to account for this error.

*Theorem 5:* Suppose that  $\hat{h}$  approximates  $h$  with an absolute error of at most  $\epsilon/n$  with probability at least  $1 - \frac{\delta}{n|\mathcal{V}|}$ . Then using  $\hat{h}$ , Algorithm 1 returns a set  $\hat{\mathcal{A}}_{\text{greedy}}$ , for which

$$F(\hat{\mathcal{A}}_{\text{greedy}}) \geq (1 - 1/e) \max_{\mathcal{A} \subseteq \mathcal{V}; |\mathcal{A}| \leq n} F(\mathcal{A}) - \epsilon \quad (10)$$

with probability at least  $1 - \delta$ .

*Proof:* Follows directly from the proof of Theorem 11 in [7]. ■

In the Monte Carlo simulation, the absolute error  $\epsilon/n$  and the confidence  $1 - \frac{\delta}{n|\mathcal{V}|}$  can be easily estimated via the standard deviation of the estimator for  $h$ . Hence, by increasing the number  $N$  of simulated runs,  $\epsilon$  and  $\delta$  in (10) can be moved arbitrarily close to zero. Note that the Hoeffding argument used by Krause and Guestrin to derive an upper bound for the error depending on the value of  $N$  cannot be extended to our case, because the random variables we consider are generally unbounded.

## V. CONTROL MODEL

In Section III we assumed the control commands  $u_{1:T}$  to be randomly chosen. However, in practice the robot often gets control commands by a human operator or operates autonomously, i.e., selects control commands depending on its current belief of the state. The control behaviors considered in this paper can be classified into the following modes:

- *Random control*, where the control commands do not depend on any other random variable (see Fig. 2 without dashed arrows).
- *External control*, where the control commands only depend on the last state of the robot, which is assumed to be known by the controller (see Fig. 2).
- *Autonomous control*, where the control commands depend on the estimated state of the robot.

In the following we show the effect of the non-random control modes to our landmark selection method.

### A. External Controls

The external control mode is modeled by the Bayes net in Fig. 2. Theorem 1 still holds for this Bayes net, because the Markov property for  $Z^{\mathcal{A}}$  conditioned on  $X$  is still applicable. Therefore, also the tight Nemhauser-bound (i.e., Theorems 3 and 4) holds. Because of the additional dependency (dashed arrows) in the Bayes net,  $p(x_{t-1} | z_{1:t-1}, u_{1:t}) = \eta p(u_t | x_{t-1}) p(x_{t-1} | z_{1:t-1}, u_{1:t-1})$ . Therefore, the factorization of the full posterior in (9) has to be extended to

$$p(x_{1:T} | z_{1:T}^{\mathcal{A}}, u_{1:T}) = \prod_{t=1}^T \eta'_t p(z_t^{\mathcal{A}} | x_t) p(x_t | x_{t-1}, u_t) p(u_t | x_{t-1}), \quad (11)$$

where  $p(u_t | x_{t-1})$  models the policy of the controller.

We assume  $p(u_t | x_{t-1})$  to be known and normally distributed, so that the entropy calculation can be easily extended to the external control mode.

### B. Autonomous Controls

In the case of fully autonomous operation of the robot, each control  $u_t$  depends on the belief of the state  $p(x_{t-1} | z_{1:t-1}, u_{1:t-1})$ , i.e., all previous control commands  $u_{1:t-1}$  and observations  $z_{1:t-1}^{\mathcal{A}}$ . Extended by these dependencies the Bayes net no longer fulfills the Markov property for  $Z^{\mathcal{A}}$  and therefore the conditional mutual information is not submodular in the autonomous mode which was a precondition for the Nemhauser approximation guarantee.

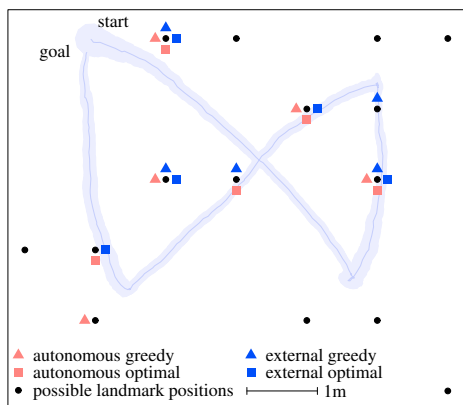


Fig. 3. The landmarks selected by our greedy algorithm and optimal landmark selections for autonomous and external controls. A sample trajectory for external controls using the greedily selected landmarks is plotted in light blue together with its  $2\sigma$  uncertainty.

Concretely, the probability that landmarks further along the trajectory become visible to the robot depends on how well the robot is localized, i.e., which landmarks it has observed before. So, adding a landmark to a small set of selected landmarks can result in a lower increase of conditional mutual information than adding the same landmark to a superset of the selected landmarks for which the probability of reaching this last landmark is higher. Such a landmark configuration obviously breaks the submodularity property. However, in the simulation experiments (see Section VI-A) the results of the greedy algorithm are still close to the optimal results even for autonomous controls.

## VI. EXPERIMENTAL RESULTS

To evaluate our approximation algorithm, we performed extensive experiments both in simulation and with a real robot. In the experiments, we consider a sensor which observes range and bearing to unique landmarks and has a limited field of view. The field of view is considered circular with radius  $2m$  in simulation, while for the real robot the rectangular field of view of its camera is used. As implementation of the Bayes filter, we use a Square-Root Unscented Kalman filter (SR-UKF) [15], which approximates the real probability distributions with  $k$ -dimensional normal distributions  $\mathcal{N}(\mu, \Sigma)$ . This allows for a straightforward calculation of the differential entropy as  $h(\mathcal{N}) = \log \sqrt{(2\pi e)^k |\Sigma|}$ .

### A. Simulation Experiments

Two different kinds of simulation experiments were performed, each with  $N = 10,000$  runs:

In the first experiment, we compare the solution produced by Algorithm 1 to the optimal solution. We selected  $n = 5$  out of  $|\mathcal{V}| = 15$  landmarks for a trajectory forming a pointed figure eight. This is the maximum problem size for which we were able to determine the optimal solution via brute force. Fig. 3 shows the landmark set  $\mathcal{A}_{\text{greedy}}$  selected by the greedy approach and the optimal landmark set  $\mathcal{A}^* = \text{argmax} F(\mathcal{A})$  both for autonomous and external controls. In both control modes, the greedy strategies place two of the five landmarks on different positions than the brute force strategies. Two




Fig. 4. The mutual information values for the four landmark sets from Fig. 3 and 95% confidence intervals of the Monte Carlo simulation.

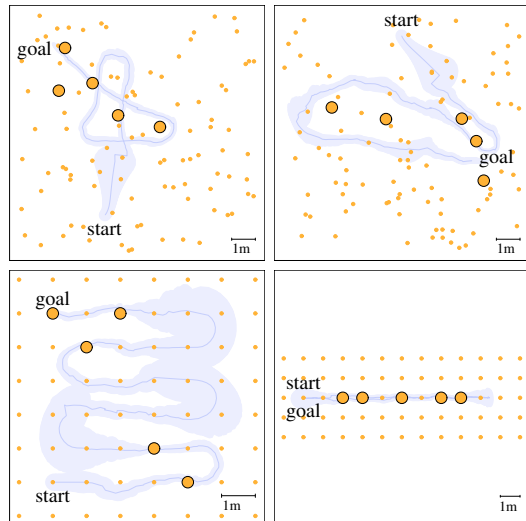


Fig. 5. Four instances in which our algorithm selected five landmarks assuming autonomous controls. The selected landmarks are shown as big orange dots, the sets  $\mathcal{V}$  as small orange dots, and one execution of the desired trajectories in blue. First row: two of the ten sampled tasks. Second row: a sweeping task and a fetch-and-return task.

landmarks are placed on the same positions in all four cases. The mutual information values  $F(\mathcal{A})$  of the four landmark sets are shown in Fig. 4. In this scenario, the approximation quality of the greedy algorithm is 99.1% of the optimal value for autonomous controls and 100.3% for external controls. The percentage higher than 100% is due to the standard deviation in the Monte Carlo simulation. So even for the case of autonomous controls, for which we did not derive a bound on the approximation error, the greedy algorithm performs well in the experiment. The absolute values in Fig. 4 are of course much lower for the external control case, as  $h(X | U)$  already is much lower due to the additional integration of  $p(u_t | x_{t-1})$  into the posterior distribution.

In the second simulation experiments, we evaluated our approach independently of a specific trajectory and a set of possible landmark positions  $\mathcal{V}$ . In order to achieve this abstraction, we randomly sampled ten trajectories and ten associated sets  $\mathcal{V}$ , each containing 100 landmarks, in a  $10m \times 10m$  environment. For each trajectory, we evaluated the selections for five and ten landmarks, so together we got 20 landmark selection tasks. Fig. 5 displays the five first landmarks our algorithm selected for two of the ten sampled tasks and the selections for two handcrafted sweeping and

TABLE I  
EVALUATION ON SAMPLED TRAJECTORIES

	mut. inf.	$h_{\text{final}}$	$\bar{d}[m]$	$d_{\text{final}}[m]$
our approach	1121.11	-4.32	0.191	0.126
uniqueness	791.11	-2.64	0.285	0.254
random	697.09	-1.70	0.348	0.409
no landmarks	0.00	2.59	1.084	1.807

TABLE II  
QUALITY IN SIMULATION AND REALITY

	mut. inf.	$h_{\text{final}}$	$\bar{d}[m]$	$d_{\text{goal}}[m]$
simulation	2278.09	-4.62	0.061	0.073
reality	2570.94	-5.33	0.068	0.040

fetch-and-return tasks. In the sampled tasks, we evaluated our approach of greedily placing landmarks to maximize mutual information against randomly placing landmarks, no landmarks, and the uniqueness maximization algorithm [8]. The latter approach maximizes the average uniqueness over all possible states  $x \in \mathcal{X}$  of the robot. It defines the uniqueness of a pose  $x$  given a map  $m$  as  $\mathcal{U}(x, m) = (\int_{\tilde{x} \in \mathcal{X}} p(z^{[x, m]} | \tilde{x}, m) d\tilde{x})^{-1}$ , where  $z^{[x, m]}$  is the maximum likelihood observation at pose  $x$  given the map  $m$ . To evaluate the landmark selections, we consider four different measures of quality: the mutual information  $I(X; Z^A | U)$ , the entropy at the final state  $h_{\text{final}} := h(X_T | Z_{1:T}^A, U_{1:T})$ , and the average and final distances  $\bar{d}$  and  $d_{\text{final}}$ , respectively, between the true pose and the mean of the belief. Table I summarizes the evaluation results for autonomous controls. Our approach based on mutual information maximization yields the best results for all four criteria. In paired sample t-tests, the differences between the results of our approach and the results of all other approaches were statistically significant at a 5% level. In fact, the highest p-value was 0.21%. For external controls, the differences in the criteria of quality are similar to the ones shown in Table I. They are also all statistically significant at a 5% level.

### B. Experiments with a Real Robot

To further validate the simulation results, we evaluated the four criteria of quality also on a real Pioneer P3-DX robot. The robot and the executed trajectory are shown in Fig. 1. We used a standard webcam pointing towards the ceiling as a sensor for detecting unique ARToolkit markers [1]. The considered set  $\mathcal{V}$  of possible landmark locations is disconnected by areas where lamps and supporting beams do not allow landmark placement (see Fig. 1). To evaluate the criteria  $\bar{d}$  and  $d_{\text{final}}$ , we obtained reference positions from laser-based Monte Carlo localization [14].

The eight landmarks depicted in Fig. 1 were selected by our algorithm. Using these landmarks, the real robot executed ten autonomous runs. Table II shows the criteria of quality for this scenario evaluated in simulation together with the average values of the ten real runs. Despite of the imperfect reference positions and remaining systematic errors the quality values in the real world experiments corresponded to the ones obtained in simulation.

## VII. CONCLUSIONS

In this paper, we presented a trajectory dependent landmark selection method for mobile robot navigation. We formulated landmark selection as the optimization problem of maximizing conditional mutual information of the states of the robot given the observations of the landmarks. This problem belongs to a class of problems which are known to be NP-hard. Using the concept of submodularity, we derived a tight constant-factor bound on the error of our polynomial time approximation algorithm. In contrast to other approaches, our algorithm is not restricted to linear systems and does not make specific assumptions on the structure of the set of possible landmark locations. Extensive experiments demonstrate that it outperforms other approaches.

## REFERENCES

- [1] M. Billinghurst and H. Kato. Collaborative augmented reality. *Communications of the ACM*, 45(7):64–70, 2002.
- [2] C.M. Bishop. *Pattern recognition and machine learning*. Springer, 1 edition, 2007.
- [3] T. Cover and J. Thomas. *Elements of Information Theory*. Wiley Interscience, 1991.
- [4] H.F. Durrant-Whyte, D. Pagac, B. Rogers, M. Stevens, and G. Nelmes. Field and service applications - an autonomous straddle carrier for movement of shipping containers. *IEEE Robotics & Automation Magazine*, 14:14–23, 2007.
- [5] J.-S. Gutmann, E. Eade, P. Fong, and M. Munich. A constant-time algorithm for vector field slam using an exactly sparse extended information filter. In *Proc. of Robotics: Science and Systems (RSS)*, 2010.
- [6] D.B. Jourdan and N. Roy. Optimal sensor placement for agent localization. *ACM Trans. Sen. Netw.*, 4(3):1–40, 2008.
- [7] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *Proc. of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2005.
- [8] D. Meyer-Delius, M. Beinhofer, A. Kleiner, and W. Burgard. Using artificial landmarks to reduce the ambiguity in the environment of a mobile robot. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2011.
- [9] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations for maximizing submodular set functions. *Mathematical Programming*, 14:265–294, 1978.
- [10] B.W. Parkinson, J.J. Spilker, P. Axelrad, and P. Enge. Global positioning system: Theory and applications. In P. Zarchan, editor, *Progress in Astronautics and Aeronautics*, volume 2. American Institute of Aeronautics and Astronautics, 1996.
- [11] P. Sala, R. Sim, A. Shokoufandeh, and S. Dickinson. Landmark selection for vision-based navigation. *IEEE Transactions on Robotics and Automation*, 22(2):334–349, 2006.
- [12] 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.
- [13] S. Thrun. Finding landmarks for mobile robot navigation. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1998.
- [14] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2006.
- [15] R. van der Merwe. *Sigma-Point Kalman Filters for Probabilistic Inference in Dynamic State-Space Models*. PhD thesis, OGI School of Science & Engineering, Oregon Health & Science University, Portland, OR, USA, 2004.
- [16] M. Vitus and C. Tomlin. Sensor placement for improved robotic navigation. In *Proc. of Robotics: Science and Systems (RSS)*, 2010.
- [17] P. Wurman, R. D’Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–19, 2008.
- [18] 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)*, 2005.