

Robust Landmark Selection for Mobile Robot Navigation

Maximilian Beinhofer

Jörg Müller

Andreas Krause

Wolfram Burgard

Abstract—Precise navigation is a key capability of autonomous mobile robots and required for many tasks including transportation or docking. To guarantee a robust and accurate localization and navigation performance, many practical approaches rely on observations of artificial landmarks. This raises the question of where to place the landmarks along the desired trajectory of the robot. In this paper, we present a novel approach to landmark selection, which aims at selecting the minimal set of landmarks that bounds the uncertainty about the deviation of the robot from its desired trajectory. At the same time the selected landmark sets are robust against the fact that a certain number of landmarks can be obscured from view during operation. Our algorithm is highly efficient due to a linearization of the whole navigation cycle and employs submodular optimization, for which strong formal bounds on the approximation quality are known. In extensive experiments, also carried out with a real robot, we demonstrate that our approach outperforms several other methods and that it enables robust autonomous robot navigation in practice.

I. INTRODUCTION

For autonomous mobile robots, being able to make distinctive observations of the environment is essential to ensure a reliable navigation performance. Due to possible ambiguities and dynamic changes, the features which are present in the environment are often insufficient for the desired accuracy in navigation. Therefore, especially in industry and logistics, many practical approaches rely on the use of artificial landmarks to achieve a robust navigation performance [8], [20]. Depending on their type, the landmarks or their placement can be expensive. Additionally, the computing power required during navigation increases with the size of the landmark set. On the other hand, landmarks can wear out over time or be obstructed by dynamic objects. Therefore, an ideal landmark placement minimizes their number, while nevertheless allowing for a certain redundancy.

In this paper, we consider the problem of selecting a set of landmarks that is suitable for robust navigation of mobile robots that repeatedly execute the same trajectory. Our approach to landmark selection builds on Bayesian A-optimal design [15], bounding the trace of the covariance of the robot's deviation throughout the whole trajectory. Thereby, our method effectively bounds the uncertainty about the deviation of the robot from its desired trajectory for all dimensions of the state space. Our approach aims at minimizing the number of landmarks needed, while still satisfying the bound on the trace even if any k of the



Fig. 1. The KARIS robot in the experimental environment. The two laser scanners on opposite corners of the robot provide a 360° field of view. The three stripes of reflective tape on the walls can be detected in the laser scans and are part of the robust landmark set selected by our approach.

selected landmarks are not observable. By choosing k , the user can trade off the number of landmarks needed against the robustness to missing landmarks.

Our approach has several characteristics that make it especially useful in practice: The robustness against missing landmarks allows the application of smooth fallback procedures: If the robot does not observe a placed landmark for some time, it can send a signal to the maintenance personnel and can still travel safely back to its parking position, leaving the workspace unobstructed for others. For landmark placement, we take into account the desired trajectory of the robot as well as its motion model and its sensor model, making the selected landmarks especially useful for the considered navigation task and the specific robot. As we linearize the model of the entire navigation cycle, our approach is highly efficient and therefore can be utilized even in large scale scenarios. We represent the space of possible landmark locations by a discrete set. Therefore, in contrast to many continuous optimization procedures, our approach does not rely on continuity assumptions for the space of possible landmark locations and is easily adjustable to new scenarios. To achieve the desired robustness in the landmark placement, our approach uses a conservative approximation of the landmark visibility that does only depend on the desired bound specified by the user. Also, we formulate our objective function in a way that allows us to utilize techniques from submodular function optimization, which come with formal approximation guarantees.

II. RELATED WORK

In the context of robot navigation, there exist several approaches to landmark selection for a robot with on-board sensors and to the similar problem of sensor location

M. Beinhofer, J. Müller and W. Burgard are with the University of Freiburg, Germany. A. Krause is with the ETH Zurich, Switzerland. This work has partly been supported by the German Research Foundation (DFG) within the Research Training Group 1103 and under contract number SFB/TR 8. beinhofe@informatik.uni-freiburg.de

selection for tracking a blind robot. For example, Sala *et al.* [16] cover an environment with landmarks so that at every position, at least $n \geq 1$ landmarks are observable. Erickson and LaValle [9] derive bounds on the maximum number of deterministically observable color-coded landmarks needed to cover a polygonal region. For a fixed set of possible robot positions, Jourdan and Roy [12] place sensors on the walls of buildings to minimize the average position error bound. Lerner *et al.* [14] select a landmark set that minimizes the weighted trace of the covariance resulting from one single observation of the landmarks. In contrast to these methods, our approach considers the full specification of the robot and its navigation task for landmark selection.

Similar to our approach, Vitus and Tomlin [18] place landmarks along a given trajectory. They use the a-posteriori covariance of the most likely run of the robot to approximate the a-priori covariance. Van den Berg *et al.* [5] consider a similar linearized system as we do, and choose one out of a set of randomly sampled sensor locations for a single sensor without observability constraints. In our previous work [3], we placed landmarks to minimize the entropy of the joint distribution of all robot states, which is a submodular optimization problem. In that approach, we applied Monte-Carlo simulations for landmark placement, which is computationally more demanding than the linearization we apply in this paper. In a second previous approach [2], we considered an efficient linearized system with a less strict observability constraint than the one presented in this paper. Therefore, that approach is not robust against missing landmarks and is highly non-submodular. Also, both our previous approaches do not consider the case of missing landmarks. See Sec. VI-B for a comparison to these methods.

To our knowledge, there exists no other approach in the landmark selection literature that takes into account the specification of the navigation task and ensures robustness against missing landmarks.

Also on the topic of submodular function optimization, there exists a large body of literature [7], [11], [13], [17], ranging from applications like Gaussian Process model fitting for ship hull inspection [11] to selecting beneficial grasping poses for robotic manipulators [17].

III. BACKGROUND AND PROBLEM STATEMENT

We consider the problem of landmark placement for a mobile robot that repeatedly and autonomously travels along the same trajectory \mathcal{T} . Hereby, the trajectory $\mathcal{T} = ((\mathbf{x}_1^*, \mathbf{u}_1^*), \dots, (\mathbf{x}_T^*, \mathbf{u}_T^*))$ is a time-discrete sequence specifying the desired robot state \mathbf{x}_t^* and control command \mathbf{u}_t^* at each time step t .

For localization, the robot has a map of the positions and the unique identities of the landmarks $\mathcal{A} = \{L_1, \dots, L_n\}$ in the area surrounding the trajectory and is equipped with a sensor to observe them. The robot takes noisy observations \mathbf{z}_t^A of the landmarks inside its actual field of view according to a sensor model $\mathbf{z}_t^A = h(\mathbf{x}_t, \text{obs}(\mathbf{x}_t, \mathcal{A}), \mathbf{n}_t)$. Here, $\text{obs}(\mathbf{x}_t, \mathcal{A})$ specifies the subset of the landmarks in \mathcal{A} which

are observable from the state \mathbf{x}_t and $\mathbf{n}_t \sim \mathcal{N}(\mathbf{0}, N_t)$ is the sensor noise, which we assume to be Gaussian distributed.

After taking an observation and updating its state estimate in a localization algorithm, the robot executes a control command \mathbf{u}_t , propagating its state \mathbf{x}_t according to a noisy motion model $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{m}_{t+1})$, where $\mathbf{m}_{t+1} \sim \mathcal{N}(\mathbf{0}, M_{t+1})$ is the motion noise. The control commands are selected by a linear-quadratic regulator (LQR) feedback controller [6]. At each time step t , the LQR controller selects the control command \mathbf{u}_t that minimizes the expected quadratic error term

$$\mathbb{E} \left[\sum_{\ell=t}^T (\Delta \mathbf{x}_\ell^T C \Delta \mathbf{x}_\ell + \Delta \mathbf{u}_\ell^T D \Delta \mathbf{u}_\ell) \right], \quad (1)$$

where $\Delta \mathbf{x}_\ell = \mathbf{x}_\ell - \mathbf{x}_\ell^*$, $\Delta \mathbf{u}_\ell = \mathbf{u}_\ell - \mathbf{u}_\ell^*$, and C and D are positive definite weight matrices.

In this closed-loop system for autonomous navigation, we consider a discrete set of possible landmark locations \mathcal{V} , and aim to select the subset $\mathcal{A} \subseteq \mathcal{V}$ of landmarks for placement that are most beneficial for the navigation task. We measure the quality of a landmark selection \mathcal{A} by considering the a-priori probability distribution of the deviation $\mathbf{x}_t - \mathbf{x}_t^*$ of the robot from its desired trajectory

$$p(\mathbf{x}_t - \mathbf{x}_t^* \mid \text{obs}(\mathbf{x}_{1:t}, \mathcal{A})) = \int \int p(\mathbf{x}_t - \mathbf{x}_t^* \mid \mathbf{u}_{1:t-1}, \mathbf{z}_{1:t}, \text{obs}(\mathbf{x}_{1:t}, \mathcal{A})) \cdot p(\mathbf{u}_{1:t-1}, \mathbf{z}_{1:t} \mid \text{obs}(\mathbf{x}_{1:t}, \mathcal{A})) d\mathbf{u}_{1:t-1} d\mathbf{z}_{1:t}. \quad (2)$$

Note that this distribution is independent of the actual values of the controls $\mathbf{u}_{1:t-1}$ and observations $\mathbf{z}_{1:t}$ that are not yet known at the moment of landmark placement. The a-priori distribution depends, however, on the number and the positions of the landmarks $\text{obs}(\mathbf{x}_{1:t}, \mathcal{A})$ that are observable during operation on the trajectory.

In this work, we consider Bayesian A-optimal design [15], i.e., we aim to select landmarks so that the trace of the estimated covariance matrix \hat{S}_t^A of $p(\mathbf{x}_t - \mathbf{x}_t^* \mid \text{obs}(\mathbf{x}_{1:t}, \mathcal{A}))$ stays below a user-defined threshold for every time step $t \in [1, T]$. Thereby, we effectively bound the uncertainty about the deviation of the robot from its desired trajectory for all dimensions of the state space.

To guarantee safe operation even if up to k landmarks are missing, we aim at finding

$$\mathcal{A}^* = \underset{\mathcal{A} \subseteq \mathcal{V}}{\text{argmin}} |\mathcal{A}| \quad (3)$$

subject to

$$\max_{\mathcal{B} \subseteq \mathcal{A}, |\mathcal{B}| \leq k} \text{tr} \left(\hat{S}_t^{\mathcal{A} \setminus \mathcal{B}} \right) \leq \epsilon_t \quad \forall t \in [1, T]. \quad (4)$$

This is the smallest set of landmarks \mathcal{A}^* that ensures a bounded trace of the estimator of the a-priori covariance. In particular, we ensure that the trace of the estimator $\hat{S}_t^{\mathcal{A} \setminus \mathcal{B}}$ of the a-priori covariance stays below ϵ_t for all $t \in [1, T]$, even if any subset $\mathcal{B} \subseteq \mathcal{A}^*$ with $|\mathcal{B}| \leq k$ is not observable during operation. Here, ϵ_t is a user-defined bound that can be set for each part of the trajectory individually. This allows

the user to specify lower ϵ_t values for critical parts of the trajectory in which a path execution with higher accuracy is required.

To efficiently place landmarks in this framework, two important issues have to be addressed: handling the combinatorial structure of the problem stated in Eqs. (3) and (4) and estimating the a-priori covariance \hat{S}_t^A . Our efficient solution to the combinatorial optimization problem is presented in the next section. In general, \hat{S}_t^A cannot be estimated in closed form, so one solution that is often applied is to approximate the high-dimensional integral defined in Eq. (2) via Monte-Carlo simulation. Monte-Carlo simulation can deal, for example, with non-linearities due to discontinuities in the observability $obs(\mathbf{x}_t, \mathcal{A})$ of landmarks depending on the actual state of the robot, but it is computationally demanding. In contrast to that, using the conservative approximation $\widehat{obs}(t, \mathcal{A})$ of the landmark observability defined in the next section, we can estimate \hat{S}_t^A efficiently by applying the linearization proposed by van den Berg *et al.* [4]. Their approach linearizes the model of the whole navigation system, consisting of observation, localization, control, and motion, resulting in a Gaussian a-priori distribution

$$\hat{p}(\mathbf{x}_t - \mathbf{x}_t^* | \widehat{obs}(1:t, \mathcal{A})) \sim \mathcal{N}(0, \hat{S}_t^A) \quad (5)$$

that can be calculated efficiently via standard matrix manipulations. This method has been applied successfully in practice [2], [5]. For completeness, it is described in detail in the appendix.

IV. EFFICIENT AND ROBUST LANDMARK SELECTION

We now describe our approach towards landmark placement. We first propose how to address the challenge of uncertain landmark observability. We then reformulate Eq. (4) such that it becomes amenable to efficient optimization.

A. Observability Constraints

For most types of landmark detection sensors (e.g., cameras, RFID readers, and laser range finders), the ability to observe a landmark L changes with the state \mathbf{x}_t of the robot due to a limited sensor range or obstacles concealing the landmark. However, in our approach it is necessary to evaluate the quality of L at a stage at which the only knowledge about \mathbf{x}_t is the condition on its covariance defined in Eq. (4), which the final landmark set shall guarantee.

Therefore, we define the approximate observability $\widehat{obs}(t, \mathcal{A})$ in a conservative way: We consider a landmark as observable at time t only if it is observable with a probability of at least $(1 - \delta)$ according to every a-priori distribution for which the bound in Eq. (4) holds. Fig. 2 shows an example of $\widehat{obs}(t, \mathcal{A})$ for a two dimensional robot state $\mathbf{x}_t = [x_t, y_t]$. Note that in the experiments, we consider a three dimensional robot state $[x_t, y_t, \theta_t]$.

We calculate $\widehat{obs}(t, \mathcal{A})$ using the following insights: As stated in Eq. (5), all estimated a-priori distributions are Gaussians. The $(1 - \delta)$ -confidence region of a Gaussian is an ellipsoid whose principal axes have the lengths $c\sqrt{\lambda_{t,i}}$. Here, $\lambda_{t,i}$ are the eigenvalues of the covariance \hat{S}_t^A and




Fig. 2. Desired robot state \mathbf{x}_t^* and landmark L . All $(1 - \delta)$ confidence regions of estimated covariances satisfying the bound in Eq. (4) (one example shown in light blue) are inside the circular region with radius $c\sqrt{\epsilon_t}$ (dark blue). According to $\widehat{obs}(t, L)$, L is visible at time t if it is inside the sensor range of every pose in the dark blue circle and the shaded area is free of obstacles that would block the robot's line of sight to the landmark.

$c = c(\delta, \dim(\mathbf{x}_t))$ is a constant that depends only on the probability δ and the dimensionality of the state space. If the bound in Eq. (4) holds, it holds that $\text{tr}(\hat{S}_t^A) \leq \epsilon_t$ and therefore $\lambda_{t,i} \leq \epsilon_t$ for all eigenvalues $\lambda_{t,i}$ of \hat{S}_t^A . Consequently, at time t , at least $(1 - \delta)$ of the probability mass of every Gaussian that satisfies Eq. (4) lies inside the sphere K with radius $c\sqrt{\epsilon_t}$ and center \mathbf{x}_t^* . If a landmark L is observable from within every state inside K , we define it to be in $\widehat{obs}(t, \mathcal{A})$, resulting in

$$\widehat{obs}(t, \mathcal{A}) = \{L \mid L \in obs(\mathbf{x}_t, \mathcal{A}) \forall \mathbf{x}_t \in K\}. \quad (6)$$

Having defined $\widehat{obs}(t, \mathcal{A})$, we can evaluate \hat{S}_t^A . Therefore, we can also evaluate the bound in Eq. (4) for a given set \mathcal{A} by iterating over all subsets of \mathcal{A} with $|\mathcal{A}| - k$ elements. Note that in the definition of $\widehat{obs}(t, \mathcal{A})$ we assume that the bound in (4) holds for \mathcal{A} . Hence, if the bound holds, then also the approximation of the observability, which was applied in the evaluation of the bound, is conservative.

B. Objective Function

Being able to evaluate Eq. (4) for a given landmark set \mathcal{A} makes it possible to run a brute force search on the power set $\mathcal{P}(\mathcal{V})$ of all possible landmark positions to find the optimal landmark set satisfying Eq. (4). However, as $\mathcal{P}(\mathcal{V})$ grows exponentially with the number of possible landmark locations $|\mathcal{V}|$, we apply an efficient approximation instead. We now show, building on techniques of Krause *et al.* [13], how the overall problem defined in Eqs. (3) and (4) can be reformulated in a way that admits highly efficient approximation algorithms.

As a first step, we define the reduction of the trace of the a-priori covariance induced by the observations of the landmarks in \mathcal{A} as

$$F_t(\mathcal{A}) = \text{tr}(\hat{S}_t^\emptyset) - \text{tr}(\hat{S}_t^A) \quad (7)$$

for every time step t . We truncate this function at the target value $\text{tr}(\hat{S}_t^\emptyset) - \epsilon_t$, leading to the function

$$F_t(\epsilon_t, \mathcal{A}) = \min \left(F_t(\mathcal{A}), \text{tr}(\hat{S}_t^\emptyset) - \epsilon_t \right). \quad (8)$$

Note that this function achieves its maximum value if and only if the target condition $tr(\hat{S}_t^A) \leq \epsilon_t$ is satisfied. We take into account the robustness against up to k missing landmarks in the objective function by considering the average over $F_t(\epsilon_t, \mathcal{A} \setminus \mathcal{B})$ for all possible subsets of missing landmarks \mathcal{B} :

$$F_t(k, \epsilon_t, \mathcal{A}) = \frac{1}{\sum_{i=0}^k \binom{|\mathcal{A}|}{i}} \sum_{\mathcal{B} \subseteq \mathcal{A}, |\mathcal{B}| \leq k} F_t(\epsilon_t, \mathcal{A} \setminus \mathcal{B}). \quad (9)$$

Similar to the function in Eq. (8), this function achieves its maximum value $tr(\hat{S}_t^{\mathcal{A}}) - \epsilon_t$ if and only if the target condition under k -robustness $tr(\hat{S}_t^{\mathcal{A} \setminus \mathcal{B}}) \leq \epsilon_t$ is satisfied for all \mathcal{B} . We finally consider multiple time steps or a whole trajectory by using the same averaging procedure, leading to

$$F(k, \epsilon_{1:T}, \mathcal{A}) = \frac{1}{T} \sum_{t=1}^T F_t(k, \epsilon_t, \mathcal{A}). \quad (10)$$

Due to its construction, this function takes on its maximum

$$c = \frac{1}{T} \sum_{t=1}^T tr(\hat{S}_t^{\mathcal{A}}) - \epsilon_t \quad (11)$$

if and only if the condition defined in Eq. (4) is satisfied. With this, we can re-formulate the problem definition stated in Eqs. (3) and (4) in terms of F as

$$\mathcal{A}^* = \underset{\mathcal{A} \subseteq \mathcal{V}}{\operatorname{argmin}} |\mathcal{A}| \quad \text{s.t.} \quad F(k, \epsilon_{1:T}, \mathcal{A}) = c. \quad (12)$$

C. Landmark Selection Algorithm

As problems of the type defined in Eq. (12) are typically NP-hard (see Sec. V for details), we apply a greedy iterative landmark selection algorithm that finds an approximate solution to (12). The procedure is stated in Algorithm 1.

Algorithm 1 Iterative approximation algorithm

Input: $\mathcal{V}, k, \epsilon_{1:T}$

Output: \mathcal{A}

$\mathcal{A} = \emptyset$

while $F(k, \epsilon_{1:T}, \mathcal{A}) < c$ **do**

$L^* = \operatorname{argmax}_{L \in \mathcal{V}} F(k, \epsilon_{1:T}, \mathcal{A} \cup \{L\})$

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

end while

return \mathcal{A}

The computation of the argmax operator in the algorithm evaluates each landmark L individually, which makes it well-suited for parallel computing.

Note that due to the usage of \widehat{obs} in the evaluation of F , which is a conservative approximation of the observability only if $F(k, \epsilon_{1:T}, \mathcal{A}) = c$, stopping the algorithm before $F(k, \epsilon_{1:T}, \mathcal{A})$ reaches c leads to landmark sets that can perform arbitrarily badly. However, for the final output set \mathcal{A} , the observability, and therefore also the condition in Eq. (4), is approximated conservatively.

D. Practical Considerations

Our algorithm can be used to guarantee a collision-free path execution with high confidence. For that, we use the same insights as for the approximation of the observability. We choose the bound ϵ_t on the trace such that the nearest static obstacle is at least $c\sqrt{\epsilon_t}$ away from the desired state \mathbf{x}_t^* for every t . To avoid collisions with moving obstacles without breaking the bound on the trace, the robot needs to stop if its path is blocked and wait until the moving obstacle left the corridor with width $c\sqrt{\epsilon_t}$ around the desired trajectory.

If the robot has to repeatedly execute the same round-trip task, our method can be adjusted to guarantee bounded traces even for a continuous long-term operation of the robot. By setting ϵ_T to at most the minimum eigenvalue of \hat{S}_1 , our algorithm produces a landmark set \mathcal{A} that guarantees that \hat{S}_T is governed by \hat{S}_1 , which enables a continuous safe operation.

V. APPROXIMATION BOUND

In this section, we provide a theoretical motivation for our approximation algorithm, which rests on the concept of submodularity, a natural diminishing returns property. For landmark sets, submodularity states that adding a landmark to an already large set of landmarks \mathcal{C} results in a smaller increase in the objective function than adding the same landmark to a subset of \mathcal{C} . Concretely, a function F is called submodular if for all $\mathcal{A} \subseteq \mathcal{C} \subseteq \mathcal{V}$ and all landmarks $L \in \mathcal{V} \setminus \mathcal{C}$

$$F(\mathcal{C} \cup \{L\}) - F(\mathcal{C}) \leq F(\mathcal{A} \cup \{L\}) - F(\mathcal{A}). \quad (13)$$

For submodular functions, problems of the type defined in Eq. (12) are called *submodular set cover* problems, and are NP-hard in general [19]. However, for these problems, Wolsey [19] showed that for greedy solutions $\mathcal{A}_{\text{greedy}}$, such as those produced by Algorithm 1, it holds that

$$|\mathcal{A}_{\text{greedy}}| \leq |\mathcal{A}^*| \left(1 + \log \max_{L \in \mathcal{V}} F(\{L\})\right), \quad (14)$$

and under natural complexity-theoretic assumptions, no efficient algorithm can provide better solutions. Hence, such greedy solutions $\mathcal{A}_{\text{greedy}}$ are near-optimal for submodular set cover problems.

Therefore, the key question is whether (or under which conditions) our objective function $F(k, \epsilon_{1:T}, \mathcal{A})$ for landmark placement is monotonic and submodular. First, note that the function $F_t(\mathcal{A})$ (Eq. (7)), from which our objective function is constructed, can be rearranged to be the sum of the variance reductions in the diagonal elements of the covariance. Das and Kempe [7] show that variance reduction, such as the summands of $F_t(\mathcal{A})$, under certain technical conditions on the covariance matrices is usually monotonic and submodular. In this case, $F(k, \epsilon_{1:T}, \mathcal{A})$ is monotonic and submodular as well, since these properties are preserved under nonnegative linear combinations and truncation [10], [13]. These observations theoretically justify our efficient approximation Algorithm 1. Similar techniques have been successfully used in other domains [11], [13], [17].




Fig. 3. A sweeping (left) and a surveillance task (right) with the 99% a-priori confidence ellipses (blue) when observing all landmarks (red triangles) in the sets that our algorithm selected for at most $k = 0$ (top), 1 (middle), or 2 (bottom) missing landmarks. The k -subsets of selected landmarks whose absence resulted in the highest simulated maximum trace are shown in pink. The red dots indicate the sets of possible landmark locations \mathcal{V} .

VI. EXPERIMENTAL RESULTS

We evaluated our approach in extensive experiments both with a simulated differential drive robot and with a real holonomic drive robot. For these robots, the state \mathbf{x}_t of the robot can be described by its pose $[x_t, y_t, \theta_t]$ in the 2d-plane. We assume that the robot is equipped with a landmark detection sensor with a circular field of view and 5 m sensor range. In the different experiments, the sensor can observe either landmarks placed on the walls or landmarks placed on the ceiling of the environment, resulting in different sets \mathcal{V} of possible landmark locations. In all experiments, we set the allowed maximum trace ϵ_t to 0.05 for all t and the probability δ in the observability constraint to 1%.

A. Evaluation of Robustness

In the first set of simulation experiments, we evaluated the robustness of our landmark sets against missing landmarks. To this end, we considered the two trajectories shown in Fig. 3, corresponding to a sweeping pattern in an obstacle-free environment and a surveillance task in an environment with obstacles. For the sweeping trajectory, our approach selected 10, 15, and 19 landmarks assuming at most zero, one, and two missing landmarks, respectively. For the surveillance task, our approach selected 12, 20, and 27 landmarks. To evaluate the effects of the linear approximation applied in the landmark selection method, we conducted Monte-

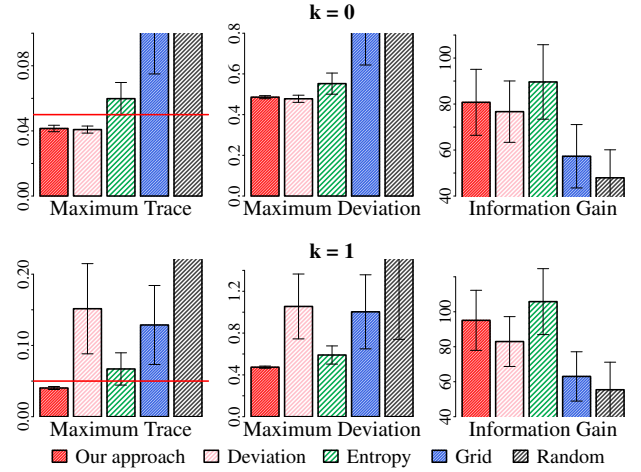


Fig. 4. Means and 95% error bars resulting from simulations on ten randomly sampled trajectories. For $k = 1$, the results for the simulations in which the one most crucial landmark was missing are shown. The horizontal red line indicates $\epsilon_t = 0.05$. The maximum deviation is stated in m .

Carlo simulations using the real, non-linear models. In the simulations, we estimated the a-priori traces $tr_{MC}(S_t)$ using the empirical distributions gained from the deviations $\mathbf{x}_t - \mathbf{x}_t^*$ observed in 1000 simulated executions of the trajectory in each scenario. For all six landmark configurations and all possible combinations of k missing landmarks, the Monte-Carlo simulations resulted in traces that were below the bound $\epsilon_t = 0.05$ for all time steps t . The maximum 0.0439 of the traces in simulation occurred in the surveillance scenario for $k = 0$.

B. Comparison to Other Approaches

In the second set of simulation experiments, we compared the landmark sets selected by our approach to the ones selected by our previous approaches [2], [3], and by two straightforward heuristics. The *Deviation* method [2] aims at minimizing the translational deviation of the robot from its trajectory in a linearized system. *Entropy* [3] minimizes the entropy of the joint distribution of $\mathbf{x}_{1:T}$ using Monte-Carlo simulations instead of linearized models. The *Grid* and *Random* heuristics place a given number of landmarks in the area observable by the robot, *Grid* in a regular grid pattern and *Random* at randomly sampled locations. For comparison, we adjusted all methods to select the same numbers of landmarks. To get scenario-independent results, we considered ten randomly chosen trajectories, each connecting six randomly sampled goal points in an area of 15 m \times 15 m. On these trajectories, our approach selected landmark sets of average sizes 5.7 and 9.3 for $k = 0$ and $k = 1$, respectively. We used Monte-Carlo simulations to evaluate the maximum trace per trajectory, the 99% quantile of the maximum translational deviation per trajectory, and the information gain of the joint distribution of $\mathbf{x}_{1:T}$. The results can be seen in Fig. 4. The *Entropy* approach results in the best information gain values, but yields suboptimal traces and deviations. This is probably due to the fact that minimizing the entropy can lead to flat but elongated covariance ellipses, which have high




Fig. 5. Pick-and-place trajectory and landmark sets selected by our approach for a constant ϵ_t of 0.05 for $k = 0$ (left) and $k = 1$ (right), and for $k = 0$ and $\epsilon_t = 0.03$ for the time steps in which the robot is inside the pick-up and deposit zones (gray areas) and 0.05 outside (middle). The trajectory goes from the start to goal 1, then to goal 2, and back to the start.

trace and deviation values. Paired sample t-tests showed that our approach results in significantly smaller maximum traces and 99% quantiles of the maximum translational deviations than *Entropy*, *Grid*, and *Random* on a 5% level for $k = 0$ as well as $k = 1$. For zero missing landmarks, the *Deviation* method produces maximum traces and deviations similar to our approach, but seems to behave unreliably in the case of a missing landmark. For $k = 1$, when the most influential landmark is hidden, it results in maximum traces and deviations that are significantly higher than the ones produced by our approach and even the *Entropy* approach on a 5% level. This is probably due to the fact that the placement algorithm employed in the *Deviation* approach aims at maximizing the distance between two consecutive landmarks, which is especially unfavorable in the case of missing landmarks, while our approach uses more sophisticated techniques from submodular function optimization to explicitly take missing landmarks into account.

Running multi-threaded on an Intel® Core™ i7 2.8GHz, the runtime of the *Entropy* approach for computing a single landmark set was 8:10h on average, while due to the linearization in landmark placement, single-threaded implementations of our approach and *Deviation* took 28:27 min and 27:22 min, respectively.

C. Landmark Selection for Changing Bounds

To demonstrate the ability of our approach to select landmarks for values of ϵ_t which vary along the trajectory, we applied our approach on the pick-and-place trajectory shown in Fig. 5. As can be seen in the figure, a higher demand for accuracy in the pick-up and deposit zones and the robustness against a missing landmark in this case lead to the same number of landmarks, but to different locations.

D. Long Term Evaluation on a Real Robot

Finally, we evaluated the landmark sets selected by our approach also on the real robot shown in Fig. 1. The robot is equipped with Mecanum wheels for omnidirectional motion and with two SICK S300 laser scanners mounted on opposite corners of the robot, providing a 360° field of view. The lasers can detect reflective markers, whose unique landmark IDs we calculated using a nearest neighbor heuristic. In a training run, we calibrated the motion noise and sensor noise of this specific robot, and used the calibrated parameters in the linearized models for landmark placement.

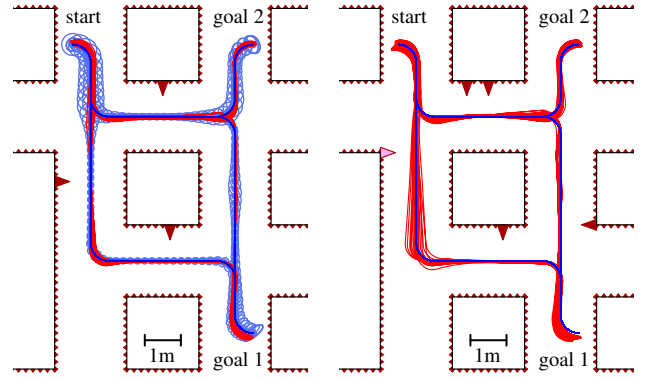


Fig. 6. Landmark sets (red triangles), desired robot path (blue) and actual robot paths (red) for $k = 0$ (left), and $k = 1$ (right). For $k = 0$, also the linearized 99% a-priori confidence regions are shown (light blue). For $k = 1$, each landmark was hidden during 10 runs of the robot. The landmark whose absence resulted in the largest deviations is marked in pink. The trajectory goes from the start to goal 1, back to the start, then to goal 2 and again back to the start.

The trajectory and the landmarks that our approach selected to ensure continuous long term operation (see Sec. IV-D) are shown in Fig. 6. To evaluate the selected landmark sets, the robot continuously executed the trajectory several times. Observing only the landmarks selected for $k = 0$, the robot autonomously executed 20 runs of the trajectory, continuously operating for one hour. The a-priori traces calculated from the deviations measured by a Motion Analysis motion capture system with ten digital Raptor-E cameras stayed considerably below 0.05, with a maximum of 0.0087 occurring close to goal 1. The measured maximum translational deviation of the robot from the desired trajectory was 0.36 m. For $k = 1$, the robot executed 50 runs of the trajectory during 2.5 hours of continuous operation. As illustrated in the video attachment, during operation each landmark was hidden from the robot during ten runs. The a-priori traces calculated from the whole dataset and also the ones calculated from each block of ten runs in which one landmark was hidden stayed below 0.05. The maximum value, 0.0444, occurred close to the lower left corner of the trajectory when the landmark marked in Fig. 6 was hidden. During the whole experiment, the measured maximum translational deviation from the desired trajectory was 0.45 m.

VII. CONCLUSIONS

In this paper, we presented a novel method to trajectory-dependent landmark placement, which is robust against missing landmarks. It keeps the trace of the a-priori covariances of all robot states below a user-defined threshold, effectively bounding the uncertainty in all dimensions of the state space. The linearized objective function in our method takes into account the full specification of the navigation task and can be evaluated efficiently. We showed how our problem can be reformulated in a way that enables the use of submodular optimization techniques, which enjoy strong theoretical guarantees. Extensive experiments, also with a real robot, demonstrate that the robustness against missing landmarks resulting from our approach is guaranteed in practice, even in

continuous long term operation. The experiments also show that our approach outperforms several other approaches to landmark selection.

APPENDIX

The Linearized System for A-Priori Covariance Estimation

In our landmark selection algorithm, we apply the efficient method for the estimation of the a-priori covariance \hat{S}_t^A introduced by van den Berg *et al.* [4], which is restated here for completeness. This method assumes that the observability of landmarks $\widehat{obs}(t, \mathcal{A})$ is independent of the concrete state \mathbf{x}_t of the robot. It uses first-order Taylor approximations to linearize the sensor model $h(\mathbf{x}_t, \widehat{obs}(t, \mathcal{A}), \mathbf{n}_t)$ and the motion model $f(\mathbf{x}_t, \mathbf{u}_t, \mathbf{m}_t)$ around the desired state \mathbf{x}_t^* and the desired control \mathbf{u}_t^* . The resulting linearized functions are

$$\mathbf{x}_t \approx f(\mathbf{x}_{t-1}^*, \mathbf{u}_{t-1}^*, \mathbf{0}) + A_t(\mathbf{x}_{t-1} - \mathbf{x}_{t-1}^*) \quad (15)$$

$$+ B_t(\mathbf{u}_{t-1} - \mathbf{u}_{t-1}^*) + V_t \mathbf{m}_t,$$

$$\mathbf{z}_t \approx h(\mathbf{x}_t^*, \widehat{obs}(t, \mathcal{A}), \mathbf{0}) + H_t(\mathbf{x}_t - \mathbf{x}_t^*) + W_t \mathbf{n}_t, \quad (16)$$

with the Jacobians

$$A_t = \frac{\partial f}{\partial \mathbf{x}}(\mathbf{x}_{t-1}^*, \mathbf{u}_{t-1}^*, \mathbf{0}), \quad B_t = \frac{\partial f}{\partial \mathbf{u}}(\mathbf{x}_{t-1}^*, \mathbf{u}_{t-1}^*, \mathbf{0}),$$

$$V_t = \frac{\partial f}{\partial \mathbf{m}}(\mathbf{x}_{t-1}^*, \mathbf{u}_{t-1}^*, \mathbf{0}), \quad (17)$$

$$H_t = \frac{\partial h}{\partial \mathbf{x}}(\mathbf{x}_t^*, \widehat{obs}(t, \mathcal{A}), \mathbf{0}), \quad W_t = \frac{\partial h}{\partial \mathbf{n}}(\mathbf{x}_t^*, \widehat{obs}(t, \mathcal{A}), \mathbf{0}).$$

In this linearized system, the Kalman filter (KF) is the minimum mean-square error estimator [1] for the a-posteriori distribution $p(\mathbf{x}_t - \mathbf{x}_t^* | \mathbf{u}_{1:t-1}, \mathbf{z}_{1:t}, \widehat{obs}(1:t, \mathcal{A}))$, estimating it as a Gaussian $\mathcal{N}(\boldsymbol{\mu}_t - \mathbf{x}_t^*, \hat{P}_t^A)$. In contrast to the extended Kalman filter (EKF) [1], we linearize around the a-priori known desired states instead of the means in the filter. Therefore, the Kalman gains K_t in the KF can be computed a-priori, without knowledge of the values of $\mathbf{u}_{1:t-1}$ and $\mathbf{z}_{1:t}$.

Applied on the mean $\boldsymbol{\mu}_t$ in the Kalman filter, the LQR controller selects control commands \mathbf{u}_t according to

$$\mathbf{u}_t - \mathbf{u}_t^* = L_t(\boldsymbol{\mu}_t - \mathbf{x}_t^*), \quad (18)$$

where L_t is the feedback matrix that minimizes the quadratic error defined in Eq. (1). With the Jacobians defined in Eq. (17), also the feedback matrices L_t can be calculated a-priori via the recursive formula $E_T = C$,

$$L_t = -(B_{t+1}^T E_{t+1} B_{t+1} + D)^{-1} B_{t+1}^T E_{t+1} A_{t+1}, \quad (19)$$

$$E_t = C + A_{t+1}^T E_{t+1} A_{t+1} + A_{t+1}^T E_{t+1} B_{t+1} L_t.$$

As can be seen in Eq. (18), the LQR controller linearly connects the Kalman filter estimate $\boldsymbol{\mu}_t$ with the true robot state \mathbf{x}_t via the selected control \mathbf{u}_t . This fact allowed van den Berg *et al.* [4] to prove that in the linear navigation system defined above, the joint a-priori distribution of the deviations of \mathbf{x}_t and $\boldsymbol{\mu}_t$ from \mathbf{x}_t^* is a Gaussian

$$\begin{bmatrix} \mathbf{x}_t - \mathbf{x}_t^* \\ \boldsymbol{\mu}_t - \mathbf{x}_t^* \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \hat{R}_t^A = \begin{bmatrix} \hat{S}_t^A & Cov(\mathbf{x}_t, \boldsymbol{\mu}_t) \\ Cov(\mathbf{x}_t, \boldsymbol{\mu}_t)^T & \hat{U}_t^A \end{bmatrix}\right),$$

whose covariance \hat{R}_t^A can be computed recursively by

$$\hat{R}_1^A = \begin{bmatrix} \hat{P}_1^A & 0 \\ 0 & 0 \end{bmatrix}, \quad \hat{R}_t^A = F_t \hat{R}_{t-1}^A F_t^T + G_t \begin{bmatrix} M_t & 0 \\ 0 & N_t \end{bmatrix} G_t^T,$$

with

$$F_t = \begin{bmatrix} A_t & B_t L_{t-1} \\ K_t H_t A_t & A_t + B_t L_{t-1} - K_t H_t A_t \end{bmatrix}, \quad (20)$$

$$G_t = \begin{bmatrix} V_t & 0 \\ K_t H_t V_t & K_t W_t \end{bmatrix}. \quad (21)$$

This recursive scheme can be computed independently of the actual values \mathbf{u}_t and \mathbf{z}_t . It efficiently estimates the a-priori covariance \hat{S}_t^A of the deviation as the upper left block of \hat{R}_t^A .

REFERENCES

- [1] Y. Bar-Shalom, T. Kirubarajan, and X. Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., 2002.
- [2] M. Beinhofer, J. Müller, and W. Burgard. Effective landmark placement for accurate and reliable mobile robot navigation. *Robotics and Autonomous Systems (RAS)*, in press, available online: <http://dx.doi.org/10.1016/j.robot.2012.08.009>.
- [3] M. Beinhofer, J. Müller, and W. Burgard. Near-optimal landmark selection for mobile robot navigation. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2011.
- [4] J. van den Berg, P. Abbeel, and K. Goldberg. LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. In *Proc. of Robotics: Science and Systems (RSS)*, 2010.
- [5] J. van den Berg, S. Patil, R. Alterovitz, P. Abbeel, and K. Goldberg. LQG-based planning, sensing, and control of steerable needles. In *Proc. of the Int. Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2010.
- [6] D. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000.
- [7] A. Das and D. Kempe. Algorithms for subset selection in linear regression. In *Proc. of the ACM Symposium on Theory of Computing (STOC)*, 2008.
- [8] H. F. Durrant-Whyte, D. Pagac, B. Rogers, M. Stevens, and G. Nemes. Field and service applications - an autonomous straddle carrier for movement of shipping containers. *IEEE Robotics & Automation Magazine*, 14:14–23, 2007.
- [9] L. H. Erickson and S. M. LaValle. Navigation among visually connected sets of partially distinguishable landmarks. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2012.
- [10] T. Fujito. Approximation algorithms for submodular set cover with applications. *IEICE Trans. on Information & Systems*, 83(3):480–487, 2000.
- [11] G. A. Hollinger, B. Englot, F. Hover, U. Mitra, and G. S. Sukhatme. Uncertainty-driven view planning for underwater inspection. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2012.
- [12] D. B. Jourdan and N. Roy. Optimal sensor placement for agent localization. *ACM Trans. on Sensor Networks*, 4(3):1–40, 2008.
- [13] A. Krause, B. McMahan, C. Guestrin, and A. Gupta. Robust submodular observation selection. *Journal of Machine Learning Research (JMLR)*, 9:2761–2801, 2008.
- [14] R. Lerner, E. Rivlin, and I. Shimshoni. Landmark selection for task-oriented navigation. *IEEE Trans. on Robotics*, 23(3):494–505, 2007.
- [15] F. Pukelsheim. *Optimal Design of Experiments*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 2006.
- [16] P. Sala, R. Sim, A. Shokoufandeh, and S. Dickinson. Landmark selection for vision-based navigation. *IEEE Trans. on Robotics and Automation*, 22(2):334–349, 2006.
- [17] J. D. Schulman, A. Lee, J. Ho, and P. Abbeel. Grasping and fixturing as submodular coverage problems. In *Proc. of the Int. Symposium of Robotics Research (ISRR)*, 2011.
- [18] M. Vitus and C. Tomlin. Sensor placement for improved robotic navigation. In *Proc. of Robotics: Science and Systems (RSS)*, 2010.
- [19] L. A. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2:385–393, 1982.
- [20] P. Wurman, R. D’Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–19, 2008.