

Two-Stage Focused Inference for Resource-Constrained Collision-Free Navigation

Beipeng Mu¹, Ali-akbar Agha-mohammadi¹, Liam Paull², Matthew Graham¹, Jonathan How¹, and John Leonard²

¹Laboratory for Information and Decision Systems

²Computer Science and Artificial Intelligence Laboratory

Massachusetts Institute of Technology, {mubp, aliapha, lpaull, mcgraham, jhow, jleonard}@mit.edu

Abstract—Long-term operations of resource-constrained robots typically require hard decisions be made about which data to process and/or retain. The question then arises of how to choose which data is most useful to keep to achieve the task at hand. As spacial scale grows, the size of the map will grow without bound, and as temporal scale grows, the number of measurements will grow without bound. In this work, we present the first known approach to tackle *both* of these issues.

The approach has two stages. First, a subset of the variables (focused variables) is selected that are most useful for a particular task. Second, a task-agnostic and principled method (focused inference) is proposed to select a subset of the measurements that maximizes the information over the focused variables. The approach is then applied to the specific task of robot navigation in an obstacle-laden environment. A landmark selection method is proposed to minimize the probability of collision and then select the set of measurements that best localizes those landmarks. It is shown that the two-stage approach outperforms both only selecting measurement and only selecting landmarks in terms of minimizing the probability of collision. The performance improvement is validated through detailed simulation and real experiments on a Pioneer robot.

I. INTRODUCTION

One of the core enabling capabilities for mobile robots operating in uncertain and GPS-denied environments is the ability to autonomously build a map and then subsequently use this map to achieve some objective. The required fidelity and choice of map representation are, in general, domain specific, and furthermore may not even be constant across a given application. Consider the following tasks:

- A robot is navigating through an uncertain environment consisting of some open areas and some narrow hallways.
- An autonomous car is operating on a road network consisting of highways and local roads.
- A marine robot is localizing a set of underwater mines amongst clutter.

In each of these cases, there is some notion of priority of landmarks in the map. In the indoor navigation scenario, it is more important to have accurate estimates of landmark locations in tight corridors to avoid collisions. Similarly on the road network, landmarks on more highly traveled roads are more useful [1] and in the underwater scenario it is more important that the robot localize the mines as opposed to the clutter. Consequently, the robot can save resources, such as

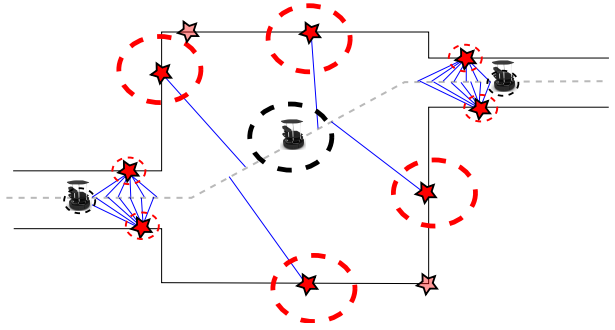


Fig. 1: Localization of some landmarks is more important than others. In the figure, it is more important to localize landmarks in narrow hallways than large rooms. As a result, far fewer measurements (blue) are required for landmarks in the open area as compared with the narrow hallway. (Landmarks are red stars, gray dashed path is trajectory of robot. Black and red dashed ellipses represent uncertainty in robot position and landmarks position respectively.

memory and computational effort, by focusing the mapping operation to more explicitly support the task.

Such complicated scenarios often require high dimensional models to represent robot poses, landmarks, and obstacles. Graphical models are a powerful tool in this case because they can explicitly represent conditional independences between variables resulting in fast inference on large scale problems [2, 3, 4]. However naively applying these methods will result in an unbounded growth in memory and computational requirements, which will ultimately result in failure for a resource-constrained system.

Recent work on map reduction either uses some criterion to discard incoming measurements [5, 6], or selectively removes nodes from the graph through marginalization followed by a sparsification procedure to maintain efficiency [7, 8, 9]. In all of these cases, the objective is map reduction while minimizing the impact on the overall quality of the map and robot trajectory.

This paper presets a more flexible map reduction framework that supports task-specific landmark prioritization. The framework is applied to the task of robot navigation and it is demonstrated that the procedure reduces the resource requirements without significantly impacting the task performance (navigation) by focusing on the important parts of the map.

The workflow is shown in Fig. 2. Assume that a prior data collection operation has been carried out. The next step is

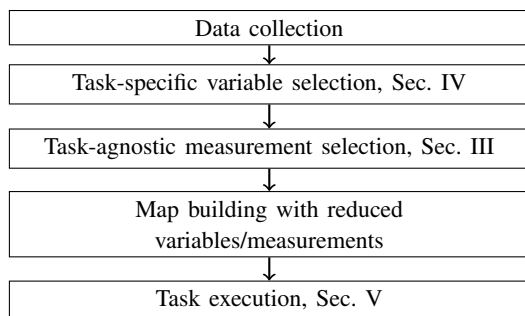


Fig. 2: Two-stage focused inference

to select the focused variables that are deemed important for task execution. Sec. IV discusses the specific case of selecting landmarks to support collision-free navigation. The third step is to select the subset of measurements that are most useful for estimating the focused variables, as described in general terms in Sec. III. The last steps are to build a map using the reduced set of variables and measurements, and then finally execute the desired task, in this case navigation. Note that in general the task may be executed many times, but the data only needs to be collected once.

In summary we claim the following contributions:

- A framework for measurement selection in the case that some variables are deemed higher priority than others (focused inference),
- A method for landmark (focused variable) selection to support the task of collision-free navigation,
- We show that focused inference over focused variables is able to minimize the probability of collision as compared to either in isolation.

II. RELATED WORK

This section reviews related work on SLAM map reduction, landmark selection and its application to autonomous navigation, and planning with uncertainty.

A. Map Reduction

Graph-based optimization approaches [10] have become popular for SLAM problems. These methods provide a naturally sparse representation of the SLAM problem that can be solved online efficiently [11]. Nevertheless, these methods do not scale constantly with time and space, and ultimately require some form of graph reduction to enable longterm operation. Such graph compression techniques can be separated into two main groups: those that incorporate each measurement and then selectively marginalize variables from the graph (SLAM back-end), and those that discard measurements directly (SLAM front-end).

In general, marginalization induces a fully connected subgraph over the Markov blanket of the marginalized variable. Most back-end map reduction approaches employ a convex optimization formulation, similar to the method introduced in [12] for sparsification, where the Kullback-Liebler divergence (KLD) between the dense subgraph and a sparse approximation is minimized subject to a consistency constraint. Carlevaris and Eustice [7] present a method called generic linear constraints that sparsifies the dense subgraph using a

Chow-Liu tree (CLT) approximation. Alternately, sparsity can be enforced through an ℓ_1 -regularization term in the KLD minimization [8], which is appealing because it does not impose a specific graph structure on the sparse approximation (e.g., a CLT). More recently, Mazuran et al. [9] improved upon previous methods by allowing non-linear measurements to approximate the dense subgraph which can be defined arbitrarily and then formulating the KLD minimization with respect to the *measurement*, rather than *state*, information matrix.

These graph reduction techniques are not concerned with selecting the nodes to be removed from the graph. Performance can degrade if the wrong landmarks are removed through marginalization since they are no longer available for subsequent loop-closures.

Front-end algorithms discard measurements before they are processed by the SLAM optimizer. Kretschmar et al. [6] propose a pose-graph compression where laser scans are selectively removed by utilizing an approximate marginalization based on a CLT, and Ila et al. [5] use an information criterion to remove uninformative loop closures.

The method proposed in this paper is a front-end approach because it directly removes measurements (either incrementally or in batch). However, the key difference as compared to [5, 6] is that measurements are selected based on how useful they are to support localization of specific landmarks, which are selected to support a specific task (e.g. collision-free navigation).

B. Landmark Selection

The notion of selecting landmarks to support localization and/or mapping has been proposed to accomplish a number of different objectives. For example, Carlone et al. [13] use a measurement selection formulation to find a maximal set of coherent measurements in the hope that this coherent set will be outlier-free. Another popular application in vision-based systems is to down-sample landmarks based on some measure of visual saliency in the hope of improving loop closure detection. Specific applications include active gaze control [14], area coverage [15], and lifelong operation of service robots [16].

More related to our motivation is the application of resource-constrained inference. For a localization and mapping objective, proposed approaches include uniform landmark selection [17] (will be referred to as “downsampling”), and entropy-based landmark selection [18].

A small number of previous works have considered resource-constrained selection of landmarks to support navigation. Strasdat et al. [19] propose a reinforcement learning based landmark selection policy to minimize the robot position error at the goal. Lerner et al. [20] consider single camera frame based landmark selection in terms of a “severity function.” Sala et al. [21] aim to choose the minimal set of landmarks such that at least k are viewable from every point in the configuration space. None of these previous works [19, 20, 21], consider the obstacles, obstacle uncertainty, or probability of collision in the landmark selection process. In contrast, our method chooses landmarks that will maximize

the probability of reaching a goal without collision, which inherently accounts for metric properties of the map, such as constrictions and tight corridors.

C. Planning with Uncertainty

This work specifically considers navigation under map and pose uncertainty as the motivating application for the focused inference framework. Much of the current literature addresses the problem of finding collision free paths in the presence of robot pose uncertainty assuming the landmark map is given. A standard measure in finding a safe path is to define a probability of collision with an obstacle. Resulting paths can be chosen that balance optimality and risk [22, 23]. In [24], an optimal path is found subject to a maximum allowable probability of collision (typically called a “chance constraint”). In [25, 26], measurement uncertainty is taken into account to compute a more accurate estimate of robot pose. The path is planned in advance assuming accurate stochastic models for motion dynamics and sensor measurements. Finally, there is a small class of planning algorithms that consider motion, measurement, and map uncertainty (e.g., [27]). However, these approaches are mainly limited to problems with small discrete state, action, or measurement spaces.

The paper presents a mapping scheme that is specifically designed to be utilized in any probabilistic navigation module. We provide a rigorous treatment of the coupling between trajectory and landmark map uncertainties, which is achieved within memory and computational constraints, making our approach applicable to low-cost robots operating with limited sensing in realistic environments.

III. MEASUREMENT SELECTION FOR INFERENCE ON FOCUSED VARIABLES

We begin, perhaps counter-intuitively, by assuming that selection of focused variables has already been performed, and then proceed to formulate the problem of measurement selection to support inference over these focused variables. This section leaves the formulation as general as possible with regards to task execution and focused variable selection, which will be discussed in the following sections.

A. Problem Formulation

Denote $X = \{X_1, \dots, X_N\}$ as a set of hidden random variables that are (partially) observable through measurements $z = \{z_1, \dots, z_K\}$ which are collected in the initial data collection phase, where K can be very large. Following the previous motivation, we also have a set of variables, the “focused” variables \tilde{X} ¹, which are a more compact representation of the variables in X : $\tilde{X} = \{\tilde{X}_1, \dots, \tilde{X}_{\tilde{N}}\}$, with $\tilde{N} \ll N$. To maintain generality, we represent the mapping from the unfocused variables to the focused variables by a *prioritization function* $w : \mathcal{R}^N \rightarrow \mathcal{R}^{\tilde{N}}$. For example, in the degenerate case of landmark selection, $\tilde{X} = w(X) = WX$ where W is an $\tilde{N} \times N$ matrix with a single 1 in each row. However, the formulation allows for more complex mappings from the unfocused to the focused set.

¹The $\tilde{\cdot}$ notation is used throughout to refer to the focused set of variables.

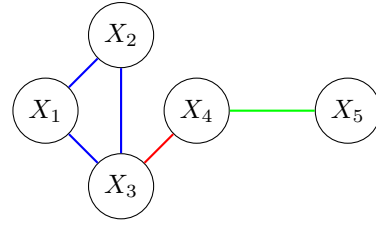


Fig. 3: Undirected graphical model. Edge colors denote maximal cliques: blue $\psi_{123}(X_1, X_2, X_3)$, red $\psi_{34}(X_3, X_4)$, and green $\psi_{45}(X_4, X_5)$.

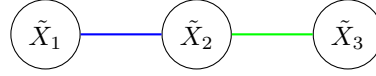


Fig. 4: Transformed undirected graphical model. New variables $\tilde{X}_1 = X_2$, $\tilde{X}_2 = \frac{1}{3}X_3 + \frac{2}{3}X_4$, $\tilde{X}_3 = X_5$. New cliques: blue $\tilde{\psi}_{12}(\tilde{X}_1, \tilde{X}_2)$, green $\tilde{\psi}_{2,3}(\tilde{X}_2, \tilde{X}_3)$.

The problem of focused measurement selection consists of choosing the best subset of the full measurement set z to optimally localize the focused variables.

Problem 1. Focused Measurement Selection: *Select the subset of measurements $z^R = \{z_1^R, \dots, z_{K^R}^R\} \subset z$, such that some information metric $f(\cdot)$ over the **focused** hidden variables \tilde{X} is maximized, subject to some cost function $g(\cdot)$ constraint on the measurement set:*

$$\begin{aligned} \max_{z^R \subset z} f(\tilde{X}; z^R) \\ \text{s.t. } g(z^R) \leq c \end{aligned} \quad (1)$$

Commonly used information metrics include entropy, mutual information, and KLD, and a common cost function is the cardinality of the set $g(z^R) = K^R$.

Graphical models, such as the Markov random field (MRF) in Fig. 3 are a compact way of representing dependencies between variables[2]. For a MRF, the joint posterior can be expressed as a product of clique potentials, $\psi_c(x_{\{c\}})$:

$$p(x|z^R) \propto \prod_{c \in \mathcal{C}} \psi_c(x_{\{c\}}) \quad (2)$$

where \mathcal{C} is the set of all cliques, $x_{\{c\}}$ are the variables in clique c . Because each clique potential $\psi_c(x_{\{c\}})$ is strictly positive, $p(X = x|z^R)$ can be written equivalently in logistic form:

$$p(x|z^R) \propto \exp\left\{\sum_{c \in \mathcal{C}} \phi_c(x_{\{c\}})\right\} \quad (3)$$

where $\phi_c(x_{\{c\}}) = \log \psi_c(x_{\{c\}})$.

The reduction from the full set of variables to the focused variables is achieved by mapping the posterior through the prioritization function $w(\cdot)$ to produce a new posterior over the focused variables:

$$p(\tilde{x}|z^R) \propto \exp\left\{\sum_{c \in \tilde{\mathcal{C}}} \tilde{\phi}_c(\tilde{x}_{\{c\}})\right\} \quad (4)$$

where $\tilde{\mathcal{C}}$ is the new (smaller) set of cliques over \tilde{X} and $\tilde{\phi}_c$ are the resulting clique potentials.

If we define the function $f(\cdot)$ in (1) to be the Shannon entropy of the conditional distribution:

$$H(\tilde{X}|z^R) = \mathbb{E}_{\tilde{X}|z^R}[-\log p(\tilde{x}|z^R)] \quad (5)$$

where \mathbb{E} is the expectation operator, then we obtain the resulting equation for the entropy of the focused variables, \tilde{X} conditional on the subset of measurements z^R as:

$$f(\tilde{X}; z^R) = H(\tilde{X}|z^R) = \mathbb{E}_{\tilde{X}|z^R}[-\sum_{c \in \tilde{C}} \tilde{\phi}_c(\tilde{x}_{\{c\}})] + C \quad (6)$$

where C is a constant.

It should be noted that computing the transformation from ϕ to $\tilde{\phi}$ may be hard in general. Furthermore, in general the graph over \tilde{X} will be much denser than the graph of X , and computing $H(\tilde{X}|z^R)$ may be computationally expensive.

However, it will be shown that $H(\tilde{X}|z^R)$ can be computed in closed form given two assumptions:

Assumption 1. *The clique potentials can be approximated as Gaussian distributions.*

Assumption 2. *The prioritization function $w(\cdot)$ is an affine transformation.*

Assumption 1 is, in fact, less limiting than the standard additive Gaussian noise assumption in the SLAM literature, since even in the case of robust cost functions we can always approximate the posterior as a Gaussian distribution using the Laplacian approximation [28]. Assumption 2 essentially requires the focused variables to be linear combinations of the original variables, which is still more general than other variable selection methods [17, 18] which restrict the set of focused variables to be a strict subset of the original set.

B. Gaussian Approximation of Clique Potentials

We begin by applying the standard method [29] of approximating a posterior over the unfocused variables (3) using a second-order Taylor series expansion of the potential functions ϕ_c at some initial guess x^* , and denote the approximated potential as $\hat{\phi}_c$:²

$$\begin{aligned} p(x|z^R) &\approx \hat{p}(x|z^R) \propto \exp\left\{\sum_{c \in \tilde{C}} \hat{\phi}_c(x_{\{c\}})\right\} \\ &= \exp\left\{\sum_{c \in \tilde{C}} \phi_c(x_{\{c\}}^*) + (x - x^*)^T \sum_{c \in \tilde{C}} \frac{\partial}{\partial x} \phi_c(x_{\{c\}}^*) \right. \\ &\quad \left. + \frac{1}{2}(x - x^*)^T \left(\sum_{c \in \tilde{C}} \frac{\partial^2}{\partial x^2} \phi_c(x_{\{c\}}^*)\right) (x - x^*)\right\}. \end{aligned} \quad (7)$$

We note that the exponential component in (7) is quadratic in x , therefore the approximation is a Gaussian distribution with information matrix, Λ_{z^R} given by the Hessian:

$$\hat{p}(x|z^R) = \mathcal{N}^{-1}(\zeta, \Lambda_{z^R}), \quad \Lambda_{z^R} = \sum_{c \in \tilde{C}} -\frac{\partial^2}{\partial x^2} \phi_c(x_{\{c\}}^*), \quad (8)$$

which can be further decomposed in the case of single and pairwise cliques as:

$$\Lambda_{z^R} = \frac{1}{2} \sum_{i,j=1..N} x_i^T \Lambda_{i,j} x_j, \quad (9)$$

where $\Lambda_{i,j} \neq 0$ only if $i = j$ or X_i and X_j are connected in the graph. We write Λ_{z^R} with subscript z^R to explicitly represent that the information is dependent on the choice of

selected measurements which will impact the structure of the graph and the resulting cliques.

One point of note is that Λ_{z^R} is inherently dependent on the linearization point chosen x^* . For computational efficiency, we pick x^* at the beginning and fix it for rest of the measurement selection procedure.

C. Affine Prioritization Function

In Sec. III-A, we defined the prioritization function $w(\cdot)$, which is a task-specific and predefined function that maps the set of all variables onto the set of focused variables. In this section we impose the restriction that this function is affine in order to provide a closed-form means of getting from (3) to (4):

$$\tilde{X} = w(X) = WX, \quad (10)$$

where $W \in \mathcal{R}^{\tilde{N} \times N}$. For example, in Fig. 4, we have:

$$\begin{bmatrix} \tilde{X}_1 \\ \tilde{X}_2 \\ \tilde{X}_3 \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1/3 & 2/3 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_W \begin{bmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \\ X_5 \end{bmatrix}. \quad (11)$$

This restriction on the prioritization to be affine guarantees that the posterior over the focused variables will still be Gaussian: $\hat{p}(\tilde{x}|z^R) = \mathcal{N}^{-1}(\tilde{\zeta}, \tilde{\Lambda}_{z^R})$. Furthermore, we can easily write an expression for the information matrix:

$$\tilde{\Lambda}_{z^R} = (W\Lambda_{z^R}^{-1}W^T)^{-1}, \quad (12)$$

as a result the approximate entropy of the focused variables given the selected measurements can be written in closed form:

$$\hat{H}(\tilde{X}|z^R) = -\frac{1}{2} \log |\tilde{\Lambda}_{z^R}| + C = \frac{1}{2} \log |W\Lambda_{z^R}^{-1}W^T| + C \quad (13)$$

We finish by restating Problem 1 based on the Gaussian approximation and restriction to affine prioritization functions:

Problem 2. Approximate Focused Measurement Selection *Select the subset of measurements $z^R = \{z_1^R, \dots, z_{K^R}^R\} \subset z$, such that approximate entropic information over the **focused** hidden variables \tilde{X} is maximized, subject to the same constraint as (1):*

$$\begin{aligned} \max_{z^R \subset z} & -\log |W\Lambda_{z^R}^{-1}W^T| \\ \text{s.t.} & g(z^R) \leq c \end{aligned}$$

D. Efficiently Solving Problem 2

Each new measurement z_k^R added to the set will introduce a new clique potential, $\phi_k(x_{\{k\}}^*)$, into the joint posterior, where the set of variables $X_{\{k\}}$ are the ones affected by measurement z_k^R . We denote the intermediate set of $k \leq K^R$ measurements that have already been selected as $z^{R_k} = \{z_1^R, \dots, z_k^R\}$.

The approximate entropy reduction (or information gain) over the focused variables brought about by adding a new measurement z_k^R is:

$$\Delta \hat{H}(\tilde{X}|z_k^R) = \hat{H}(\tilde{X}|z^{R_{k-1}}) - \hat{H}(\tilde{X}|z^{R_k}), \quad (14)$$

is the value that we want to maximize. In the following theorem we show that this quantity can be efficiently computed:

²The $\hat{\cdot}$ notation is used throughout to refer to the Gaussian approximation.

Theorem 1. *The approximate reduction in entropy over the focused variables brought about by introducing the new measurement z_k^R will be:*

$$\Delta \hat{H}(\tilde{X}|z_k^R) = -\frac{1}{2} \log \left| I - (I + J_k^T \Lambda_{z^{R_{k-1}}}^{-1} J_k)^{-1} L_k^T \tilde{\Lambda}_{z^{R_{k-1}}} L_k \right| \quad (15)$$

where $L_k \triangleq W \Lambda_{z^{R_{k-1}}}^{-1} J_k$, J_k is the measurement covariance weighted stacked Jacobian [8], and W is the affine prioritization function.

Proof: We proceed similarly to [5], but with the added complication that there is a transformation from the unfocused to the focused variables ($\tilde{X} = WX$) required to evaluate (14).

From (8), the information matrix after the introduction of z_k^R will be:

$$\Lambda_{z^{R_k}} = \Lambda_{z^{R_{k-1}}} - \frac{\partial^2 \phi_k(x_{\{k\}}^*)}{\partial x^2} = \Lambda_{z^{R_{k-1}}} + J_k J_k^T \quad (16)$$

where $J_k \in \mathcal{R}^{N \times |x_{\{k\}}|}$, with $|x_{\{k\}}|$ the number of variables in clique k and only blocks corresponding to the variables in the clique $x_{\{k\}}$ being non-zero. The new entropy after adding measurement z_k^R can be evaluated as

$$\hat{H}(\tilde{X}|z_k^R) = \frac{1}{2} \log |W(\Lambda_{z^{R_k}})^{-1} W^T| \quad (17)$$

Then the reduction in entropy, $\Delta \hat{H}(\tilde{X}|z_k^R)$, is

$$\begin{aligned} & -\frac{1}{2} \log |W \Lambda_{z^{R_k}}^{-1} W^T| + \frac{1}{2} \log |W \Lambda_{z^{R_{k-1}}}^{-1} W^T| \\ &= -\frac{1}{2} \log \frac{|W(\Lambda_{z^{R_{k-1}}} + J_k J_k^T)^{-1} W^T|}{|(W \Lambda_{z^{R_{k-1}}}^{-1} W^T)|}. \end{aligned} \quad (18)$$

By applying the matrix inversion lemma:

$$(A + BB^T)^{-1} = A^{-1} - A^{-1}B(I + B^T A^{-1}B)^{-1}B^T A^{-1},$$

and the determinant property $|I + AB| = |I + BA|$, (18) can be reduced to:

$$-\frac{1}{2} \log \left| I - (I + J_k^T \Lambda_{z^{R_{k-1}}}^{-1} J_k)^{-1} L_k^T \Lambda_{z^{R_{k-1}}} L_k \right|,$$

which is the required result. \blacksquare

1) *Computational Complexity of (15):* Similar to [5], we have avoided the need to compute the entropy over the entire variable set. However, unlike in [5] where the calculation of information gain scales only with the size of the measurements, we have a slightly more complicated scenario because of the prioritization transformation W . Upon further inspection of (15), calculating $J_k^T \Lambda_{z^{R_{k-1}}}^{-1} J_k$ has a computational complexity of $\mathcal{O}(|x_{\{k\}}|^4)$ (similar to [5]). In addition, note that calculating the i th element of L_k (which is computed as $L_k^i = W^i \Lambda_{z^{R_{k-1}}}^{-1} J_k$, where W^i is i th row of W) requires us to check the submatrix of $\Lambda_{z^{R_{k-1}}}^{-1}$ that corresponds to the non-zero elements of W^i and J_k . Therefore, the overall complexity of computing L_k is $\mathcal{O}(\|W\|_0 |x_{\{k\}}| + |x_{\{k\}}|^4)$, where $\|W\|_0$ is the number of non-zero elements in the prioritization transformation. Typically the clique size, $|x_{\{k\}}| \ll N$, and since $\tilde{N} \ll N$ we should expect that $\|W\|_0 \ll N$, therefore the overall complexity of computing L_k is much less than the problem size N . Algorithm 1 summarizes the measurement selection approach.

Algorithm 1 Measurement Selection for Prioritized Landmarks

Input: Initial information matrix $\Lambda_{z^{R_0}}$, focused variables \tilde{X} , all measurements z , budget c

Output: easurements z^R

- 1: $k \leftarrow 0$, $z^{R_0} \leftarrow \emptyset$
 - 2: **while** $g(z^{R_k}) < c$ **do**
 - 3: $k \leftarrow k + 1$
 - 4: $z_k^* = \operatorname{argmax}_{z_k^R \in z \setminus z^{R_{k-1}}} \Delta \hat{H}(\tilde{X}|z_k^R)$
 - 5: $z^{R_k} = z^{R_{k-1}} \cup \{z_k^*\}$
 - 6: $\Lambda_{z^{R_k}} = \Lambda_{z^{R_{k-1}}} + J_k J_k^T$
 - 7: $\tilde{\Lambda}_{z^{R_k}} = (W \Lambda_{z^{R_k}}^{-1} W^T)^{-1}$
 - 8: **end while**
-

IV. FOCUSED MAPPING FOR NAVIGATION

In this section, we return to the question of how to select the focused variables that are an input to the measurement selection scheme described in the previous section. The selection of focused variables mainly depends on the robot's tasks and goals. Here we specifically consider the case of collision-free autonomous navigation.

The robot uses a map of landmarks to localize itself and navigate through an environment. To build the initial map, assume that a robot with image and depth/laser sensors (referred to as the data gathering (DG) robot) has been manually operated in the environment once to gather the initial dataset. The DG-robot's trajectory can be represented as a sequence of random variables $X = \{X_1, \dots, X_T\}$. In GPS-denied environments X is not directly observable. However, the robot can always measure the incremental change between two sequential poses (odometry), for example from an IMU or wheel encoder.

There also exists a set of landmarks from which focused landmarks can be selected. Denote the set of landmarks as $L = \{L_1, L_2, \dots, L_N\}$.

A. Selection of Focused Landmarks

While it could be possible to have thousands of landmarks, often a small, carefully chosen subset can lead to sufficiently accurate navigation. In particular, for resource-constrained systems, reducing the number of landmarks will significantly reduce the computation required for data association, which will in turn enable faster and more efficient on-line trajectory planning and navigation.

Narrow passages are challenging for collision-free autonomous robot motion planning. However, in the case of high map and robot uncertainty, the "narrowness" should be redefined. For example, a "geometrically" wide passage might still be problematic for a robot that does not have access to accurate landmark information and thus has poor localization accuracy. We refer to this passage as being "geometrically wide" but "probabilistically narrow" and we will formalize these terms below.

The evaluation of probabilistic narrowness involves two key components: 1) Calculating an estimate of the robot's position uncertainty and 2) Calculating the probability of collision based on the robot's uncertainty and the distance to obstacles as determined by the DG-robot.

1) *Robot position uncertainty*: To generate a path-independent estimate of the robot's position uncertainty at any given point x along the robot's path, we use the concept of belief stabilization introduced in [30]. Assume we have a closed-loop controller that can stabilize the system to x in which the landmarks viewable during stabilization are fixed (this assumption can be relaxed, c.f. [31]). Such a controller is typically comprised of an estimator and a separated controller. The estimator generates an *a posteriori* distribution over all robot poses based on the existing map of landmarks and the local observation of landmarks. Given these estimates, the separated controller will generate a control signal that drives the robot toward x . To design an analytic measure of narrowness we rely on a simple Linear Quadratic Gaussian (LQG) controller which is comprised of a Kalman filter as an estimator and a linear quadratic regulator as a separated controller. It can be shown that starting from any $\Sigma_0 > \Sigma^*(x)$, estimation covariance decreases monotonically and approaches the covariance $\Sigma^*(x)$, which is the fixed point of a Riccati recursion at location x :

$$\{\Sigma^* = Q + A(\Sigma^* - \Sigma^* H^T (H \Sigma^* H^T + R)^{-1} H \Sigma^*) A^T\}_x \quad (19)$$

The Jacobians, A and H are computed by linearizing the measurement model at point x on the path, where Q and R are the process and measurement noise respectively. We associate with each pose x a set of visible landmarks $L(x)$. The value of the steady-state covariance will be dependent on the set of visible landmarks as expressed through the measurement noise covariance R . Note that Jacobian matrices and set of visible landmarks would be different for different points on the path. The main computational advantage in using this measure is that $\Sigma^*(x)$ only depends on x not the path that leads to x .

2) *Collision probability*: Using $\mathcal{N}(x_t, \Sigma^*(x_t))$ as a path-independent measure of uncertainty for each point x_t on the path, the collision probability can be computed by a Monte Carlo method: for each x_t sample the normal distribution and count the number of samples that fall into the obstacle region. However, we propose an alternate method that utilizes an approximate measure and is computationally cheaper. We simply approximate the probability of collision, P_c , as the Mahalanobis distance between x_t and the closest obstacle point. In other words, P_c measures how many standard deviations obstacles are away from the mean of the distribution over x_t :

$$P_c(x_t) = (x_t - x_t^{obst})^T (\Sigma^*(x_t))^{-1} (x_t - x_t^{obst}) \quad (20)$$

where the relative vector $x_t - x_t^{obst}$ is the smallest reading received from the laser range finder in its local frame.

The landmark selection problem is framed as finding the poses along the trajectory with the highest collision probability then select landmarks that can reduce this pose's uncertainty.

Problem 3. Minimum Collision Probability Landmark Selection: Select the α landmarks $L_f \subset L$, such that the maximum probability of collision is minimized:

$$\begin{aligned} \min_{L_f \subset L} \max_{x_t} P_c(x_t) \\ \text{s.t. } |L_f| \leq \alpha \end{aligned} \quad (21)$$

Algorithm 2 Minimum Collision Probability Landmark Selection

Input: Robot poses X , landmarks L , odometry noise Q , measurement noise R , budget α

Output: Selected landmarks L_k^f

- 1: $k \leftarrow 0, L_0^f \leftarrow \emptyset$
 - 2: Compute $P_c(x_t)$ for all $x_t \in X$ by (19) and (20)
 - 3: $P(l) = \max_{x_t \text{ s.t. } l \in L(x_t)} P_c(x_t)$ for all $l \in L$
 - 4: **while** $k < \alpha$ **do**
 - 5: $k \leftarrow k + 1$
 - 6: $l_k^* \leftarrow \operatorname{argmax}_{l \in L \setminus L_{k-1}^f} P(l)$
 - 7: $L_k^f = L_{k-1}^f \cup \{l_k^*\}$
 - 8: update $P_c(x_t)$ for x_t that can observe l_k^*
 - 9: update any $P(l)$ that may have changed
 - 10: **end while**
-

The problem is solved by greedily selecting landmarks. At each iteration, pick the landmark l^* that is associated with maximal $P_c(x_t)$, then update $P_c(x_t)$ for all the x that can observe l^* . The algorithm is summarized in Alg. 2. Since the number of poses that can observe any individual landmark is low, this greedy approximation can be computed efficiently.

B. Focused Measurement Selection for Navigation

Here we detail how to apply this landmark selection scheme resulting from solving Problem 3 to the measurement selection process described in Sec. III. Recall that the joint log probability of an MRF is proportional to the sum of its potentials $\phi(X_c, L_c | o, z)$:

$$p(X, L) \propto \exp\left(\sum_{c \in \mathcal{C}} \phi_c(X_c, L_c | o, z)\right). \quad (22)$$

An odometry measurement, \tilde{o} , adds a potential between two subsequent poses $\phi_t(X_t, X_{t+1}; o_t)$ and a landmark measurement, z , adds a potential between a pose and a landmark $\phi_t^i(X_t, L_i; z_t^i)$.

The focused variables are selected focused landmarks: $\tilde{X} = L_f \subset L$. Denote the rest of unfocused landmarks and robot poses as $\tilde{X}' = \{L \setminus L_f, X\}$. then the affine prioritization function represented by W is

$$\tilde{X} = W \begin{bmatrix} L_f \\ \tilde{X}' \end{bmatrix}, \quad W = \begin{bmatrix} I_{\tilde{N} \times \tilde{N}} & \mathbf{0}_{\tilde{N} \times N - \tilde{N}} \end{bmatrix}. \quad (23)$$

If we apply the standard assumption that the odometry and landmark measurement are corrupted by additive Gaussian noise [10], then the potentials are:

$$\begin{aligned} X_0 &\sim \mathcal{N}(\mu_0, \Sigma_0) & X_{t+1} &= f(X_t, o_t) + \eta, \quad \eta \sim \mathcal{N}(0, R_t) \\ \phi_t(x_t, x_{t+1}) &= -\frac{1}{2} (x_{t+1} - f(x_t, o_t))^T R_t^{-1} (x_{t+1} - f(x_t, o_t)) \\ z_t^i &= g(x_t, l_i) + \nu, \quad \nu \sim \mathcal{N}(0, Q_{t,i}) \\ \phi_t^i(x_t, l_i) &= -\frac{1}{2} (z_t^i - g(x_t, l_i))^T Q_{t,i}^{-1} (z_t^i - g(x_t, l_i)) \end{aligned}$$

Now the W matrix and the potentials can be used in Alg. 1 for the second stage: focused measurement selection.

V. NAVIGATION

The final step is to actually use the reduced map to navigate with a resource constrained robot.

The built map is represented by a set of stochastic landmarks $L \sim \mathcal{N}(\hat{L}, R^L)$. Denote the desired trajectory to T steps as $(x_t^d)_{t=0}^T$. Assume the robot starts at $x_t \sim \mathcal{N}(\hat{x}_0, \Sigma_0)$. Represent the control problem as:

$$\begin{aligned} \min_{u_t(\cdot)} J &= \mathbb{E} \left[\sum_{t=0}^T (x_t - x_t^d)^T W_x (x_t - x_t^d) + u_t^T W_u u_t \right] \\ x_{t+1} &= f(x_t, u_t) + w_t, \quad w_t \sim \mathcal{N}(0, Q_t) \\ z_t &= h(x_t, L) + v_t \quad v_t \sim \mathcal{N}(0, R_t), \quad L \sim \mathcal{N}(\hat{L}, R^L) \end{aligned} \quad (24)$$

where W_x and W_u represent the cost for robot states and control inputs. Q_t and R_t represent process and measurement noise respectively. The system is first linearized at the nominal trajectory $(x_t^d)_{t=0}^T$:

$$\begin{aligned} x_{t+1} &= x_{t+1}^d + A_t(x_t - x_t^d) + B_t u_t + w_t \\ z_t &= h(x_t^d, \hat{L}) + H_t(x_t - x_t^d) + M_t^L h(L - \hat{L}) + v_t \end{aligned} \quad (25)$$

where A_t , B_t , H_t and M_t^L represent the Jacobian of $f()$ and $g()$ at the linearization point. The optimization of the linearized system can be obtained by a Kalman filter and a LQR controller. The equivalent measurement noise covariance for the Kalman filter is:

$$\bar{R}_t = (M_t^L)^T R^L M_t^L + R_t \quad (26)$$

which incorporates both the landmark uncertainty as $(M_t^L)^T R^L M_t^L$ and measurement noise R_t . The LQR controller will have the following form:

$$u_t = -(B_t^T S_t B_t + W_u)^{-1} B_t^T S_{k+1} A_t (\hat{x}_t^+ - x_t^d) \quad (27)$$

where \hat{x}_t^+ the estimated mean at the t -th time step and S_t is obtained by solving a discrete algebraic Riccati equation.

A more general and detailed description of navigation is presented in the supplemental material.

VI. EXPERIMENTS

Both simulation and real world experiments were done to test the proposed algorithms.

A. Simulation

Fig. 5 shows the simulation environment. We performed repeated trials in the simulation environment to test the ability of our proposed scheme to be able to navigate from a start point without collision. In total there are 74 landmarks (black stars), which the robot can observe with a range/bearing measurement if landmark is in the robot's field of view. We assume that our resource constrained robot has a budget of $\alpha = 30$ landmarks (Alg. 2) and $K^R = 90$ measurements (Alg. 1). We compare six cases: (1) *optimal*: all landmarks and all measurements; (2) *focus, all*: focused landmarks and all measurements; (3) *focus, info*: focused landmarks and measurements selected based on information gain (our proposed method); (4) *focus, DS*: focused landmarks and uniform downsampling of measurements; (5) *full, info*: all landmarks and measurements selected based on information gain; (6) *full, DS*: all landmarks with uniform downsampling of measurements. Case (1) violates both stated budgets and should be seen as a lower bound on possible performance. Case (2) violates the K^R budget but gives a lower bound on the attainable localization accuracy

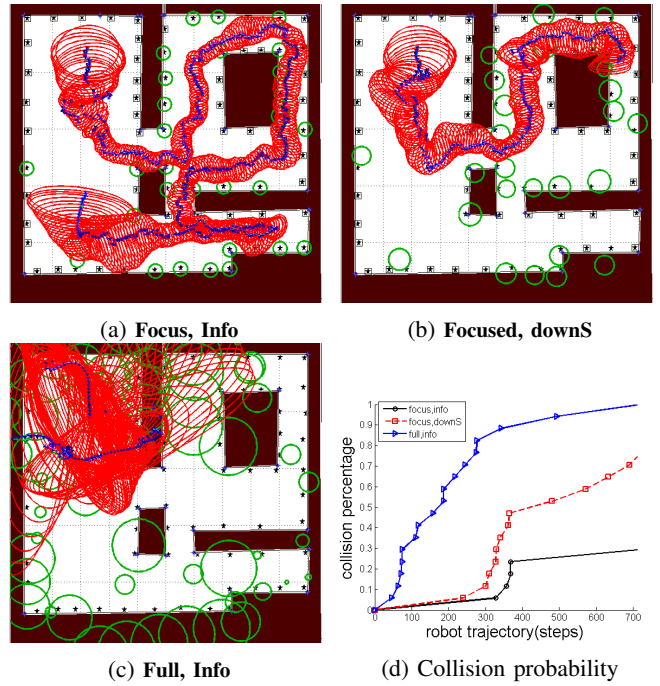


Fig. 5: Navigation with focused landmark selection and map building. Green circles represent selected landmarks with their size representing uncertainty. Blue lines are the nominal trajectories each robot wants to follow with red circles representing pose uncertainty. The focused approach has much less uncertainty in narrow passages, thus lower collision probability compared to the unfocused case

of the focused set. Cases (5) and (6) violates the α budget but, as we show, actually results in much poorer overall performance since the measurement budget is still respected and consequently all landmarks are localized indiscriminately resulting in unnecessary landmark estimate precision in wide corridors where it is not needed.

Fig. 5 (a), (b), and (c) shows three sample trials for case (3), (4) and (5) respectively and (d) shows the overall probabilities of collision obtained from all trials. The trials are stopped whenever there is an actual collision with an obstacle. The unfocused landmarks are distinguished by small black boxes around them. The 3σ -ellipses corresponding to landmark uncertainty are shown in green. The blue line represents the nominal trajectory that the robot is trying to follow. The red ellipses represent the uncertainties of the robot along the nominal trajectory. In the focused landmark selection cases, the proposed procedure picks the landmarks that contribute more in reducing robot's uncertainty in desired regions (narrow passages) and spend the computational budget to reduce the uncertainty of these focused landmarks. In the *full, info* case, the measurements are spread across landmarks. As a result, each landmark gets very little resource thus the method failed to recover a meaningful map for navigation, thus showing that, in this case, measurement selection alone would not produce an acceptable result. In the *focus, downS* case, more resources are spent on focused landmarks, thus the map is more accurate than *full, info*, but measurement selection is not based on how much they contribute to uncertainty reduction, thus the landmark positions are much less accurate than *focused, info*. The key point to note about the proposed approach Fig. 5-

TABLE I: Comparison of simulated mapping results

	No. observed lms	No. focused lms	Entropy on focused lms	Error on all observed lms	Error on focused lms
optimal	74	30	-61.3362	0.0204	0.0206
focus, all	30	30	-54.7506	0.0303	0.0303
focus, info	30	30	7.9911	0.3872	0.3872
focus, downS	30	26	66.6130	2.8026	2.8026
full, info	74	30	39.9598	22.3963	25.5736
full, downS	44	19	100.9545	2.0227	2.0456

(a) is that the robot uncertainty is preferentially reduced in the areas of the environment where the corridors are tight and there is a higher chance of collision.

We further compare all six cases from a mapping perspective in Table I based on three metrics: the entropy of the focused landmarks, the error of all the observed landmarks, and the error of the focused landmarks (which will be the same as the previous in all cases where only focused landmarks are being estimate). From the proposed focused two-step landmark and measurement selection approach achieves the best performance.

B. Hardware experiment

In the real-world experiment, we ran a Pioneer robot in a cluttered office space. The robot is equipped with an RGB camera and a 270 degree field of view laser range finder. Figure 6 shows the floor plan of the environment. AprilTags [32] were put up to create an initial pool of landmarks. The odometry measurements are obtained by matching consecutive laser scans using the open source canonical scan matcher [33]. The landmark measurements are obtained by running the AprilTag detector with the RGB images, which gives the relative orientation and range of the tags in the robot’s frame[32]. A summary of the dataset is in Table II. The selected measurements and odometry information are then fed into a standard SLAM solver g2o [34] to optimize the graph. It is important to note that we did not use the laser for the full SLAM solution, only for generating reliable odometry measurements.

TABLE II: Office Dataset

length	8min
# odometry measurements	1547
# landmark measurements	1039
# landmarks	76

30 landmarks were selected as focused based on the collision probability and 90 measurements are selected to optimize the graph. Fig. 7 compares the entropy on the 6 cases mentioned in simulation. The proposed focused mapping has significant lower uncertainties and is closest to bounds given by case (1) and (2).

Fig. 8 compares mapping results of case (1) (*optimal*), case (3) (*focus, info*), case (5) (*full, info*) and case (6) (*full, downS*). The rebuilt robot trajectory is shown with a color map, where the red color on the trajectory indicates the risky (close to obstacles) regions and blue indicates the safer regions. Magenta circles represent landmarks with the size representing its uncertainty. The focused approach (Fig. 8b) can concentrate the measurements on the narrow passage and door way, resulting in less uncertainty there. The other approaches scatter the measurements across different landmarks, and thus have much higher landmark uncertainty in narrow passages.




Fig. 6: Floor plan

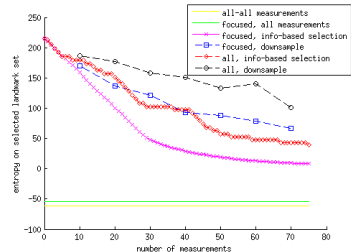


Fig. 7: Information gained

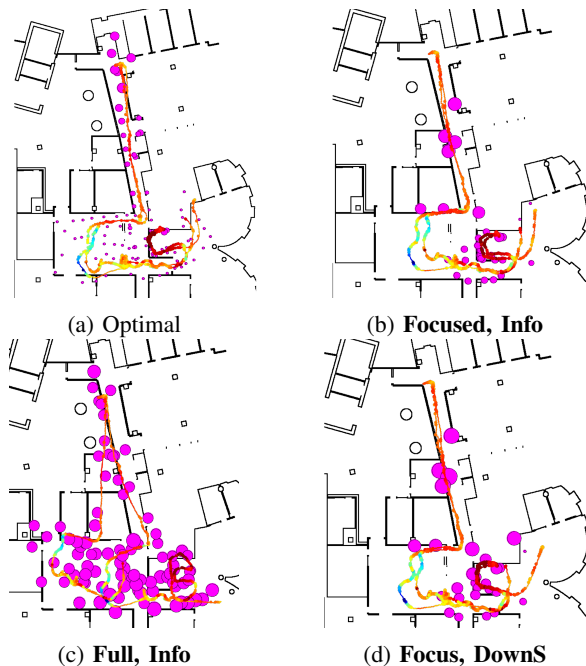


Fig. 8: Mapping results. Color line represent robot’s risk of collision. Magenta circles represent landmarks with the size representing its uncertainty. The proposed two-stage approach (Focus,Info) outperforms either measurement reduction (Full, Info) or landmark reduction (Focus, DownS) isolated

VII. CONCLUSION

This paper presented a two-stage landmark and measurement selection procedure for resource-constrained robots operating in unknown or uncertain environments. In the first stage, a set of focused variables is selected that are most important for accomplishing a specific task. The second stage, measurements are selected to maximize the information gain on these focused variables. The two-stage procedure is then applied to the task of collision-free robot navigation in an obstacle-laden environment. Simulation and hardware results demonstrated that the approach can identify a relevant subset of landmarks and accurately localize them to reduce the probability of colliding with obstacles as compared with only landmark selection or measurement selection in isolation.

ACKNOWLEDGMENTS

This research is supported in part by ARO MURI grant W911NF-11-1-0391, ONR grant N00014-11-1-0688 and NSF Award IIS-1318392.

REFERENCES

- [1] Ted J. Steiner, Guoquan Huang, and John J. Leonard. Location utility-based map reduction. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, May 2015.
- [2] Martin J. Wainwright and Michael I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning*, 1(1-2):1–305, January 2008.
- [3] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1st edition, 2007.
- [5] V. Ila, J.M. Porta, and J. Andrade-Cetto. Information-based compact pose SLAM. *Robotics, IEEE Transactions on*, 26(1):78–93, Feb 2010.
- [6] Henrik Kretzschmar and Cyrill Stachniss. Information-theoretic compression of pose graphs for laser-based SLAM. *The International Journal of Robotics Research*, 31(11):1219–1230, 2012.
- [7] N. Carlevaris-Bianco and R.M. Eustice. Generic factor-based node marginalization and edge sparsification for pose-graph SLAM. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 5748–5755, May 2013.
- [8] Guoquan Huang, M. Kaess, and J.J. Leonard. Consistent sparsification for graph optimization. In *Mobile Robots (ECMR), 2013 European Conference on*, pages 150–157, Sept 2013.
- [9] Mladen Mazuran, Gian Diego Tipaldi, Luciano Spinello, and Wolfram Burgard. Nonlinear graph sparsification for SLAM. In *Robotics Science and Systems Conference*, pages 1–8, 2014.
- [10] Frank Dellaert and Michael Kaess. Square root SAM: Simultaneous location and mapping via square root information smoothing. *International Journal of Robotics Research*, 25(12):1181–1203, 2006.
- [11] M. Kaess, A. Ranganathan, and F. Dellaert. iSAM: Incremental smoothing and mapping. *Robotics, IEEE Transactions on*, 24(6):1365–1378, Dec. 2008.
- [12] John Vial, Hugh Durrant-Whyte, and Tim Bailey. Conservative sparsification for efficient and consistent approximate estimation. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, pages 886–893, Sept 2011.
- [13] Luca Carlone, Andrea Censi, and Frank Dellaert. Selecting good measurements via l_1 -relaxation: A convex approach for robust estimation over graphs. In *Intelligent Robots and Systems (IROS), 2014 IEEE/RSJ International Conference on*, pages 2667–2674. IEEE, 2014.
- [14] S. Frintrop and P. Jensfelt. Attentional landmarks and active gaze control for visual SLAM. *Robotics, IEEE Transactions on*, 24(5):1054–1065, Oct 2008.
- [15] Ayoung Kim and Ryan M. Eustice. Active visual slam for robotic area coverage: Theory and experiment. *The International Journal of Robotics Research*, 2014.
- [16] S. Hochdorfer and C. Schlegel. Landmark rating and selection according to localization coverage: Addressing the challenge of lifelong operation of slam in service robots. In *Intelligent Robots and Systems (IROS), 2009. IEEE/RSJ International Conference on*, pages 382–387, Oct 2009.
- [17] G. Dissanayake, H. Durrant-Whyte, and Tim Bailey. A computationally efficient solution to the simultaneous localisation and map building (slam) problem. In *Robotics and Automation (ICRA), 2000. IEEE International Conference on*, volume 2, pages 1009–1014 vol.2, 2000.
- [18] Sen Zhang, Lihua Xie, and M.D. Adams. Entropy based feature selection scheme for real time simultaneous localization and map building. In *Intelligent Robots and Systems (IROS). 2005 IEEE/RSJ International Conference on*, pages 1175–1180, Aug 2005.
- [19] H. Strasdat, C. Stachniss, and W. Burgard. Which landmark is useful? learning selection policies for navigation in unknown environments. In *Robotics and Automation (ICRA), 2009 IEEE International Conference on*, pages 1410–1415, May 2009.
- [20] R. Lerner, E. Rivlin, and I. Shimshoni. Landmark selection for task-oriented navigation. *Robotics, IEEE Transactions on*, 23(3):494–505, June 2007.
- [21] P. Sala, R. Sim, A. Shokoufandeh, and S. Dickinson. Landmark selection for vision-based navigation. *Robotics, IEEE Transactions on*, 22(2):334–349, April 2006.
- [22] Yifeng Huang and K. Gupta. Collision-probability constrained PRM for a manipulator with base pose uncertainty. In *Intelligent Robots and Systems (IROS), 2009 IEEE/RSJ International Conference on*, pages 1426–1432, oct. 2009.
- [23] Brandon D. Luders. *Robust Sampling-based Motion Planning for Autonomous Vehicles in Uncertain Environments*. PhD thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, May 2014.
- [24] L. Blackmore, M. Ono, and B.C. Williams. Chance-constrained optimal path planning with obstacles. *Robotics, IEEE Transactions on*, 27(6):1080–1094, dec. 2011.
- [25] Samuel Prentice and Nicholas Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *The International Journal of Robotics Research*, 2009.
- [26] Jur van den Berg, Pieter Abbeel, and Kenneth Y. Goldberg. LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. *The International Journal of Robotics Research*, 30(7):895–913, 2011.
- [27] Hanna Kurniawati, Tirthankar Bandyopadhyay, and Nicholas M Patrikalakis. Global motion planning under uncertain motion, sensing, and environment map. *Autonomous Robots*, 33(3):255–272, 2012.
- [28] D.M. Rosen, Guoquan Huang, and J.J. Leonard. Inference over heterogeneous finite-/infinite-dimensional systems using factor graphs and gaussian processes. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 1261–1268, May 2014.
- [29] David J. C. MacKay. *Information Theory, Inference & Learning Algorithms*. Cambridge University Press, New York, NY, USA, 2002.
- [30] Ali-akbar Agha-mohammadi, Suman Chakravorty, and Nancy Amato. FIRM: Sampling-based feedback motion planning under motion uncertainty and imperfect measurements. *International Journal of Robotics Research (IJRR)*, 33(2):268–304, 2014.
- [31] Anastasios I. Mourikis and Stergios I. Roumeliotis. Predicting the performance of cooperative simultaneous localization and mapping (C-SLAM). *The International Journal of Robotics Research*, 25(12):1273–1286, 2006.
- [32] Fast odometry from vision. <http://april.eecs.umich.edu/wiki/index.php/AprilTags>.
- [33] Laser scan tools. http://wiki.ros.org/scan_tools.
- [34] g2o: A general framework for graph optimization. <https://openslam.org/g2o.html>.