

Virtual Maps for Autonomous Exploration of Cluttered Underwater Environments

Jinkun Wang, Fanfei Chen, Yewei Huang, John McConnell, Tixiao Shan and Brendan Englot

Abstract—This paper presents a novel exploration framework for underwater robots operating in cluttered environments, built upon simultaneous localization and mapping (SLAM) with imaging sonar. The proposed system comprises path generation, place recognition forecasting, belief propagation and utility evaluation using a *virtual map*, which estimates the uncertainty associated with map cells throughout a robot’s workspace. We evaluate the performance of this framework in simulated experiments, showing that our algorithm maintains a high coverage rate during exploration while also maintaining low mapping and localization error. The real-world applicability of our framework is also demonstrated on an underwater remotely operated vehicle (ROV) exploring a harbor environment.

I. INTRODUCTION

Simultaneous localization and mapping (SLAM) has been well studied in theory and applied successfully on real sensing platforms for state estimation and map-building using data collected passively [1]. However, it’s still a challenge for an autonomous vehicle to actively map an unknown environment, properly managing the trade-off between exploration speed and state estimation quality. Improving the capability of autonomous exploration is beneficial for many robot mapping tasks, especially in scenarios where teleoperation is limited or infeasible due to constrained communication, e.g., in unknown subsea environments with underwater robots.

The autonomous exploration problem is generally solved in three stages: path generation, utility evaluation and execution. First, we identify candidate waypoints or generate a sequence of actions to follow, which is typically achieved by enumerating frontiers or by employing sampling-based path planning methods. The selected path is usually straightforward to execute using feedback controllers, thus leaving us with a fundamental problem of designing a utility function to measure path optimality. Essentially, it should capture the exploration-exploitation dilemma, i.e., a balancing of visiting unknown areas to reduce map uncertainty, and revisiting known areas to seek better localization and map accuracy.

Predicting the impact of future actions on system uncertainty remains an open problem [1], and this is particularly

This work was supported by a grant from Schlumberger Technology Corporation, and by NSF Grants IIS-1652064 and IIS-1723996.

F. Chen, Y. Huang, J. McConnell, and B. Englot are with the Department of Mechanical Engineering, Stevens Institute of Technology, Hoboken, NJ, USA (e-mail: {fchen7, yhuang85, jmconn1, benglot}@stevens.edu).

J. Wang was previously with Stevens Institute and is now with Amazon Robotics, Louisville, CO, USA (email: jwang92@stevens.edu).

T. Shan was previously with Stevens Institute of Technology and is now with SRI International, Princeton, NJ, USA (email: tixiao.shan@sri.com).

true for unobserved landmarks or other environmental features that may provide useful relative measurements. Our previous work on expectation-maximization (EM) exploration [2] introduced the concept of a *virtual map* composed of virtual landmarks acting as proxy for a real feature-based map, on which we are able to predict the uncertainty resulting from future sensing actions. Since every virtual landmark is connected with robot poses that can observe it, the metrics for exploration and localization are unified as the determinant of the virtual landmarks’ error covariance matrix.

In this present paper, we extend the EM exploration algorithms employed in past works [2], [3] to the real-world application that originally inspired this work - the exploration of cluttered underwater environments by a robot equipped with a multibeam imaging sonar. In doing so, this paper provides the following contributions:

- A more detailed exposition of processes for belief propagation (over candidate robot actions, and over virtual landmarks) that support the efficient implementation of EM exploration in real-time, at large scales;
- A thorough evaluation of our framework’s performance in simulated environments, examining the tradeoffs between a robot’s rate of exploration, localization and map accuracy in comparison with other algorithms;
- The first instance, to our knowledge, of fully autonomous exploration of a real-world, obstacle-filled outdoor environment, in which an underwater robot directly incorporates its SLAM process and predictions based on that process into its decision-making.

II. RELATED WORK

The various factors considered in our mobile robot exploration problem, localization uncertainty, map uncertainty, and coverage rate, have been considered individually and in combination across a large body of prior work [4], [5], [6], [7], [8], [9]. A body of relatively recent work has considered autonomous exploration and active perception with underwater robots in particular. This work has emphasized mapping with sonar, which can perceive obstacles at relatively long ranges in all types of water. Vidal et al. [10] first employed sonar-based occupancy mapping, in 2D, to explore cluttered underwater environments with an autonomous underwater vehicle (AUV), which repeatedly selects nearby viewpoints at the map’s frontiers. Subsequent work by Palomeras et al. [11] extended the approach to 3D sonar-based mapping and exploration of cluttered underwater environments using

a next-best-view framework that weighs travel distance, frontiers, and contour-following when selecting a new viewpoint. In addition to view planning, active SLAM has been applied to underwater robots, used to incorporate planned detours into visual inspections that achieve loop closures [12], and to deviate from three-dimensional sonar-based next-best-view planning to achieve a loop closure when localization uncertainty exceeds a designated threshold [13]. Of these two works, the former navigates at a fixed standoff from a ship hull's surface, and the latter explores and maps an indoor tank environment. In the sections that follow, we will describe our proposed EM exploration framework and its application to underwater mapping with sonar, demonstrating that it achieves far lower localization and mapping errors than frontier-based [14] and next-best-view [8] exploration algorithms, and that the localization and mapping errors realized are comparable to that of [13], while achieving a superior rate of exploration.

III. EM EXPLORATION

We address the problem of autonomous exploration for a range-sensing mobile robot in an initially unknown environment. Our robot performs SLAM and constructs an occupancy grid map as it explores. We assume a bounded 2D workspace $V \subset \mathbb{R}^2$ where all discretized cells \mathbf{m}_i are initialized as *unknown* $P(m_i = 1) = 0.5$. A *frontier* is defined as the boundary where free space meets unmapped space. The exploration is considered complete if no frontier can be detected. However, highly uncertain poses are likely to result in complete, yet inaccurate occupancy grid maps, limiting the usefulness of information gained by exploring unknown space. Assuming the environment contains individual landmarks $L = \{\mathbf{l}_k\}$, apart from discovering more landmarks, minimizing the estimation error is equally crucial.

A. Simultaneous Localization and Mapping

We use a *smoothing*-based approach rather than a *filtering*-based approach, adopting incremental smoothing and mapping [15] to repeatedly estimate the entire robot trajectory. This affords us the flexibility to adopt either landmark-based SLAM, or pose SLAM, scalably over long-duration missions. Importantly, it also allows corrections to be made throughout the mission to both the robot's estimated pose history, and to the map of the environment resulting from that pose history.

Let a mobile robot's motion model be defined as

$$\mathbf{x}_i = f_i(\mathbf{x}_{i-1}, \mathbf{u}_i) + \mathbf{w}_i, \quad \mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, Q_i), \quad (1)$$

and let the measurement model be defined as

$$\mathbf{z}_{ij} = g_{ij}(\mathbf{x}_i, \mathbf{l}_j) + \mathbf{v}_{ij}, \quad \mathbf{v}_{ij} \sim \mathcal{N}(\mathbf{0}, R_{ij}), \quad (2)$$

where we assume the data association between $\mathbf{x}_i, \mathbf{l}_j$ is known.

Given measurements $\mathcal{Z} = \{\mathbf{z}_k\}$, we can obtain the best estimate of the entire trajectory $\mathcal{X} = \{\mathbf{x}_i\}$ and observed landmarks $\mathcal{L} = \{\mathbf{l}_j\}$,

$$\mathcal{X}^*, \mathcal{L}^* | \mathcal{Z} = \underset{\mathcal{X}, \mathcal{L}}{\operatorname{argmin}} P(\mathcal{X}, \mathcal{L} | \mathcal{Z}). \quad (3)$$

The maximum a posteriori (MAP) estimate can be used by maximizing the joint probability, which afterwards leads to a nonlinear least-squares problem. By constructing a graph representation and linearizing nonlinear functions, the marginal distributions and joint marginal distributions, both of which are Gaussian, can be extracted using graphical model-based inference algorithms [15].

B. EM-Exploration

In the formulation of the SLAM problem as a *belief net* [16], the solution is obtained by maximizing the joint probability distribution,

$$\mathcal{X}^*, \mathcal{L}^* = \underset{\mathcal{X}, \mathcal{L}}{\operatorname{argmax}} \log P(\mathcal{X}, \mathcal{L}, \mathcal{Z}), \quad (4)$$

where $\mathcal{X}, \mathcal{L}, \mathcal{Z}$ are robot poses, landmarks, and measurements respectively. During exploration, we are confronted with unknown landmarks that haven't been observed yet. Therefore, we introduce the concept of *virtual landmarks* \mathcal{V} as latent variables, which describe potential landmark positions that would be observed when following the planned path. Then the objective is to maximize the following marginal model,

$$\begin{aligned} \mathcal{X}^* &= \underset{\mathcal{X}}{\operatorname{argmax}} \log P(\mathcal{X}, \mathcal{Z}) \\ &= \underset{\mathcal{X}}{\operatorname{argmax}} \log \sum_{\mathcal{V}} P(\mathcal{X}, \mathcal{Z}, \mathcal{V}). \end{aligned} \quad (5)$$

The above equation involves unobserved variables, which can be approached intuitively using an expectation-maximization (EM) algorithm as follows,

$$\text{E-step: } q(\mathcal{V}) = p(\mathcal{V} | \mathcal{X}^{\text{old}}, \mathcal{Z}) \quad (6)$$

$$\text{M-step: } \mathcal{X}^{\text{new}} = \underset{\mathcal{X}}{\operatorname{argmax}} \mathbb{E}_{q(\mathcal{V})} [\log P(\mathcal{X}, \mathcal{V}, \mathcal{Z})]. \quad (7)$$

In the E-step, latent virtual landmarks are computed based on the current estimate of the trajectory and the history of measurements. In the M-step, a new trajectory is selected such that the expected value of joint probability, given the virtual landmark distributions, is maximized. The iterative algorithm alternates between the E-step and M-step, but each iteration is accomplished by the execution of actions and the collection of measurements.

The equation above poses a challenge for efficient solution due to the exponential growth of potential virtual landmark configurations with respect to the number of virtual landmarks. Inspired by classification EM algorithms, an alternative solution would add a classification step (C-step) before the M-step to provide the maximum posterior probability estimate of the virtual landmark distributions,

$$\text{C-step: } \mathcal{V}^* = \underset{\mathcal{V}}{\operatorname{argmax}} p(\mathcal{V} | \mathcal{X}^{\text{old}}, \mathcal{Z}) \quad (8)$$

$$\text{M-step: } \mathcal{X}^{\text{new}} = \underset{\mathcal{X}}{\operatorname{argmax}} \log P(\mathcal{X}, \mathcal{V}^*, \mathcal{Z}). \quad (9)$$

If we further assume measurements are assigned to maximize the likelihood

$$\mathcal{Z} = \underset{\mathcal{Z}}{\operatorname{argmax}} h(\mathcal{X}, \mathcal{V}),$$

then the joint distribution can be expressed as a multivariate Gaussian centered at the proposed poses and landmark positions, and the covariance can be approximated by the information matrix inverse,

$$P(\mathcal{X}, \mathcal{V}, \mathcal{Z}) \sim \mathcal{N}\left(\begin{bmatrix} \mathcal{X} \\ \mathcal{V} \end{bmatrix}, \begin{bmatrix} \Sigma_{\mathcal{X}\mathcal{X}} & \Sigma_{\mathcal{X}\mathcal{V}} \\ \Sigma_{\mathcal{V}\mathcal{X}} & \Sigma_{\mathcal{V}\mathcal{V}} \end{bmatrix}\right). \quad (10)$$

The solution of Eq. (9) is equivalent to evaluating the log-determinant of the covariance matrix,

$$\operatorname{argmax}_{\mathcal{X}} \log P(\mathcal{X}, \mathcal{V}^*, \mathcal{Z}) = \operatorname{argmin}_{\mathcal{X}} \log \det(\Sigma). \quad (11)$$

This implies that the performance metric for our proposed exploration is consistent with the D-optimality criterion in active SLAM [17], except that the subjects considered include unobserved landmarks.

C. Belief Propagation on Candidate Actions

Since we are more interested in the uncertainty of the virtual landmarks and the most recent pose \mathbf{x}_{T+N} at step T with planning horizon N , we can marginalize out irrelevant poses in $\Sigma_{\mathcal{X}\mathcal{X}}$, ending up with $\Sigma_{\mathbf{x}_{T+N}}$. Typically, there exist thousands of virtual landmarks, thus approximation of $\Sigma_{\mathcal{V}\mathcal{V}}$ is critical for real-time applications. Combined with pose simplification, we can obtain that, for a positive definite covariance matrix,

$$\log \det(\Sigma) < \log \det(\Sigma_{\mathbf{x}_{T+N}}) + \sum_k \log \det(\Sigma_{\mathbf{v}_k}), \quad (12)$$

where $\Sigma_{\mathbf{v}_k}$ is the diagonal block involving the k th virtual landmark in $\Sigma_{\mathcal{V}\mathcal{V}}$. This approximation is reasonable considering an overestimate of information using the introduced virtual landmarks. We will discuss the process for estimating $\Sigma_{\mathbf{x}_{T+N}}$ in this section and $\Sigma_{\mathbf{v}_k}$ in the next section.

The SLAM problem defined in Section III-A can be approximated by performing linearization and solved by iteration through the linear form given by

$$\delta^* = \operatorname{argmin}_{\delta} \frac{1}{2} \|A\delta - \mathbf{b}\|^2, \quad (13)$$

where A represents the Jacobian matrix and the vector \mathbf{b} represents measurement residuals. The incremental update δ in the above equation is obtained by solving the system,

$$A^\top A \delta = A^\top \mathbf{b}. \quad (14)$$

It can also be formulated as $\Lambda \delta = \boldsymbol{\eta}$ where $\Lambda = A^\top A$ is known as the information matrix. In general, the covariance matrix may be obtained by inverting the information matrix

$$\Sigma = \Lambda^{-1} = (A^\top A)^{-1}. \quad (15)$$

As shown in [18], the recovery of block-diagonal entries corresponding to pose covariances can be implemented efficiently. As a result, we concern ourselves with the problem of updating uncertainty upon the arrival new measurements assuming that $\Sigma_{\mathbf{x}_i}$ are given.

Covariance recovery is performed in three steps. First, we compute diagonal entries in the covariance matrix. Second, we ignore loop closure measurements and propagate the

covariance using only odometry measurements. This open-loop covariance recovery is given by

$$\Sigma_{\mathbf{x}_{i,k+t}} = \Sigma_{\mathbf{x}_{i,k+t-1}} F_{k+t}^\top, \quad (16)$$

where $F_{k+t} = \frac{\partial \mathbf{f}_{k+t}}{\partial \mathbf{x}_{k+t}}$ is the Jacobian matrix of \mathbf{f}_{k+t} with respect to pose $\mathbf{x}_{i,k-1}$. The equation is applied recursively and the initial value $\Sigma_{\mathbf{x}_{i,k}}$ can be computed by

$$\Lambda \Sigma_{\mathbf{x}_{i,k}} = I_k, \quad R \Sigma_{\mathbf{x}_{i,k}} = R^\top I_k, \quad (17)$$

where $\Sigma_{\mathbf{x}_{i,k}}$ represents the cross-covariance between past poses and the current pose and I_k is a sparse block column matrix with an identity block only at the position corresponding to pose k . The solution of Equation (17) is obtained by Cholesky decomposition on the right (the R is available immediately after incremental update in iSAM2 [15]), whose computational complexity, for sparse R with N_{nz} nonzero elements, is $O(N_{nz})$.

Finally, we use the Woodbury formula to update the covariance matrix [19],

$$\Sigma' = \Sigma + \Delta \Sigma, \quad \Delta \Sigma = -\Sigma A_u^\top (I + A_u \Sigma A_u^\top)^{-1} A_u \Sigma, \quad (18)$$

where A_u is a Jacobian matrix with each block row corresponding to one loop closure measurement. Using the above formula results in a highly efficient update, as it avoids the inversion of a large dense matrix $(A^\top A + A_u^\top A_u)^{-1}$. In total, our belief propagation process requires the recovery of block columns for which the respective pose appears in the measurements, and a matrix inversion of a relatively small matrix with the number of block rows equal to the number of loop closure measurements.

D. Belief Propagation on Virtual Landmarks

Let us assume the robot following a certain path is able to take measurements from landmarks in the surrounding environment. Let $\{\mathbf{x}_i \in \mathbf{SE}(2)\}$ be the robot poses that observe the same landmark $\mathbf{l} \in \mathbb{R}^2$. In the following derivation we do not distinguish between virtual landmark \mathbf{v} and actual landmark \mathbf{l} (the process will accommodate both). The measurement \mathbf{z}_i can be obtained from the following sensor model

$$\mathbf{z}_i = \tilde{\mathbf{z}}_i + \mathbf{v}_i = h_i(\mathbf{x}_i, \mathbf{l}) + \mathbf{v}_i, \quad \mathbf{v}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_i). \quad (19)$$

We further assume that the sensor model is invertible, i.e., that we are able to predict landmark positions given robot poses and their corresponding measurements. Mathematically, the Jacobian matrix has full rank, $\operatorname{rank}(\frac{\partial h_i}{\partial \mathbf{l}}) = 2$. One example of such a model is bearing-range measurements, which are commonly produced by sonars and laser range-finders. In the following, we will use the inverse sensor model for convenience,

$$\mathbf{l} = h_i^{-1}(\mathbf{x}_i, \tilde{\mathbf{z}}_i) = h_i^{-1}(\mathbf{x}_i, \mathbf{z}_i - \mathbf{v}_i), \quad \mathbf{v}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_i). \quad (20)$$

The Jacobian matrices of the inverse sensor model are represented as $\mathbf{H} = \frac{\partial h_i^{-1}}{\partial \mathbf{x}}$, $\mathbf{G} = \frac{\partial h_i^{-1}}{\partial \mathbf{z}}$. Given the covariance matrices of poses $\{\Sigma_i | \Sigma_i \succ 0\}$, we intend to provide a consistent estimate of a landmark's covariance without the computation of $\{\Sigma_{ij} | i \neq j\}$.




Fig. 1: Covariance intersection across several robot poses is used to compute an upper bound (red ellipses) on the actual landmark covariance (green ellipses). Raw range-bearing returns from the object represented by landmark l , matched with their nearest neighbors across consecutive robot poses using iterative closest point (ICP), are indicated as blue dots.

Suppose we obtain two measurements of the same landmark at two distinct poses $\mathbf{x}_i, \mathbf{x}_j$, resulting in a joint distribution

$$\begin{bmatrix} \mathbf{l}_i \\ \mathbf{l}_j \end{bmatrix} = \begin{bmatrix} h^{-1}(\mathbf{x}_i, \mathbf{z}_i) \\ h^{-1}(\mathbf{x}_j, \mathbf{z}_j) \end{bmatrix}. \quad (21)$$

$$\begin{aligned} \text{cov} \left(\begin{bmatrix} \mathbf{l}_i \\ \mathbf{l}_j \end{bmatrix} \right) &= \begin{bmatrix} \Sigma_i^{\mathbf{l}} & \Sigma_{ij}^{\mathbf{l}} \\ \Sigma_{ij}^{\mathbf{l}\top} & \Sigma_j^{\mathbf{l}} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{H}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_j \end{bmatrix} \begin{bmatrix} \Sigma_i & \Sigma_{ij} \\ \Sigma_{ij}^{\top} & \Sigma_j \end{bmatrix} \begin{bmatrix} \mathbf{H}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{H}_j \end{bmatrix}^{\top} \\ &\quad + \begin{bmatrix} \mathbf{G}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_j \end{bmatrix} \begin{bmatrix} \mathbf{R}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{R}_j \end{bmatrix} \begin{bmatrix} \mathbf{G}_i & \mathbf{0} \\ \mathbf{0} & \mathbf{G}_j \end{bmatrix}^{\top} \end{aligned}$$

Let $\mathbf{P}_{i1} = \mathbf{H}_i \Sigma_i \mathbf{H}_i^{\top} \succ \mathbf{0}$, $\mathbf{P}_{i2} = \mathbf{G}_i \mathbf{R}_i \mathbf{G}_i^{\top} \succ \mathbf{0}$, $\mathbf{P}_{ij} = \mathbf{H}_i \Sigma_{ij} \mathbf{H}_j^{\top}$.

$$\text{cov} \left(\begin{bmatrix} \mathbf{l}_i \\ \mathbf{l}_j \end{bmatrix} \right) = \begin{bmatrix} \Sigma_i^{\mathbf{l}} & \Sigma_{ij}^{\mathbf{l}} \\ \Sigma_{ij}^{\mathbf{l}\top} & \Sigma_j^{\mathbf{l}} \end{bmatrix} = \begin{bmatrix} \mathbf{P}_{i1} + \mathbf{P}_{i2} & \mathbf{P}_{ij} \\ \mathbf{P}_{ij}^{\top} & \mathbf{P}_{j1} + \mathbf{P}_{j2} \end{bmatrix}$$

We obtain two covariance estimates independently from two poses and we will use split covariance intersection (SCI) [20] to compute an upper bound on the actual landmark covariance as follows

$$\hat{\Sigma}^{\mathbf{l}, -1} = \left(\frac{1}{\omega} \mathbf{P}_{i1} + \mathbf{P}_{i2} \right)^{-1} + \left(\frac{1}{1-\omega} \mathbf{P}_{j1} + \mathbf{P}_{j2} \right)^{-1}, \omega \in [0, 1], \quad (22)$$

where ω can be optimized via

$$\omega^* = \underset{\omega \in (0,1)}{\text{argmin}} \det(\hat{\Sigma}^{\mathbf{l}}). \quad (23)$$

Readers are referred to [20] for analysis of SCI and proof of the upper bound. But the core idea is that if we construct an optimal linear, unbiased estimator $\hat{\mathbf{l}} = \mathbf{K}_i \mathbf{l}_i + \mathbf{K}_j \mathbf{l}_j$ ($\mathbf{K}_i + \mathbf{K}_j = \mathbf{I}$), then it can be proved that $\hat{\Sigma}^{\mathbf{l}} \succeq \mathbb{E}[(\hat{\mathbf{l}} - \mathbf{l})(\hat{\mathbf{l}} - \mathbf{l})^{\top}]$. Therefore we are able to further approximate Eq. (12) by

$$\log \det(\Sigma) < \log \det(\Sigma_{\mathbf{x}_{T+N}}) + \sum_k \log \det(\hat{\Sigma}_{\mathbf{v}_k}). \quad (24)$$

We demonstrate the process of incremental split covariance intersection in Fig. 1. At each step, a landmark covariance estimate derived from the measurement collected at a specific pose is indicated as a dark blue ellipse. After the first step, we are able to fuse landmark observations from different poses via covariance intersection, using Equation (22), as shown by red ellipses. It is evident that the resulting ellipse from SCI (red) contains the true result obtained from SLAM (green).

E. Extensibility to Pose SLAM

From the above derivation, an upper bound on the uncertainty of a real or virtual landmark may be computed from multiple poses that observe it. However it is worth investigating the meaning of this process in the context of pose SLAM, where no landmarks are incorporated into our optimization. To simplify the problem, we assume environmental features are measured and the transformation between poses with overlapping observations is derived using iterative closest point (ICP) [21]. We visualize such a scenario using hypothetical range-bearing observations of an object (which may be alternately represented as a landmark) in the leftmost plot in Figure 1. In this example, the observations are normally distributed and centered at the actual landmark location. We assume this set of range-bearing returns will be matched with those from other poses using ICP. In ICP, a point is associated to its nearest neighbor and the matched feature is given by

$$\tilde{\mathbf{l}} = \underset{\mathbf{l}_i}{\text{argmin}} \|\mathbf{l}_i - \mathbf{l}\|_2. \quad (25)$$

It is subsequently trivial to show that

$$\Sigma^{\mathbf{l}} \succeq \mathbb{E}[(\hat{\mathbf{l}} - \mathbf{l})(\hat{\mathbf{l}} - \mathbf{l})^{\top}] \succeq \mathbb{E}[(\tilde{\mathbf{l}} - \mathbf{l})(\tilde{\mathbf{l}} - \mathbf{l})^{\top}]. \quad (26)$$

The distribution of $\tilde{\mathbf{l}}$ is visualized at every step in Figure 1. While we aim to select exploration actions that reduce the (virtual) landmark covariances in landmark-based SLAM, we essentially minimize the closeness of measurements in ICP-based pose SLAM. Because association error is one major cause of ICP error, tightly clustered target points greatly contribute to registration performance.

F. Virtual Map

How can we predict unobserved landmarks without prior knowledge of the characteristics of an environment? We can approach this question by making a conservative assumption that any location which hasn't been mapped yet has a virtual landmark. Additionally, the probability that a location is potentially occupied with a landmark can be captured using an occupancy grid map. Let $P(m_i \in \mathcal{M})$ denote the occupancy of a discretized cell, then we define a virtual map \mathcal{V} consisting of virtual landmarks with probability

$$P(v_i = 1) = \begin{cases} 1, & \text{if } P(m_i = 1) \geq 0.5 \\ 0, & \text{otherwise.} \end{cases} \quad (27)$$

In its definition, existing landmarks have been incorporated into the virtual map, which is essential because minimizing the uncertainty of observed landmarks is also desired. An example of an occupancy map used to maintain virtual landmarks is provided in Figure 2, where it is used for belief propagation in support of planning with an underwater robot.

Although we can use the same occupancy grid map from path planning, which is typically high-resolution, to construct the virtual map using Eq. (27), it has been demonstrated in our earlier work [2] that a low-resolution virtual map provides similar exploration performance but requires less computation. Therefore, in our experiments, we downsample the occupancy grid map from 0.2m resolution to $d = 2$ m, shown in Fig. 2. The update of the virtual map and estimation of virtual landmark covariances are described in Alg. 1.

Algorithm 1 Virtual Landmark Covariance Estimate

Require: virtual map resolution d , trajectory \mathcal{X} and measurements \mathcal{Z} , candidate path $X_{T:T+N}$ and future measurements $Z_{T:T+N}$

- 1: $\mathcal{M} \leftarrow \text{UpdateOccupancyMap}(\mathcal{X}, \mathcal{Z})$
- 2: $\mathcal{V} \leftarrow \text{Downsample}(\mathcal{M}, d)$ ▷ Eq. (27)
- 3: $\Sigma' \leftarrow \text{UdateCovarianceDiagonal}(\mathcal{X}, \mathcal{Z}, X_{T:T+N}, Z_{T:T+N})$
▷ Eq. (16) – (18)
- 4: **for** $\mathbf{v}_k \in \{\mathbf{v}_i \in \mathcal{V} \mid P(\mathbf{v}_i = 1) = 1\}$ **do**
- 5: $\Sigma_{\mathbf{v}_k} \leftarrow \text{CovarianceIntersect}(\Sigma', X_{T:T+N}, \mathbf{v}_k)$ ▷ Eq. (19) – (23)
- 6: **end for**

G. Motion Planning

To carry out the M-step of our algorithm introduced in Equation (9), given the distribution of virtual landmarks, path candidates must be generated and evaluated using the selected utility function. If we are to consider global paths that will be followed over a long span of time, we must take into account two types of actions, exploration and place-revisiting [5]. *Exploration actions* normally have destinations near frontier locations where mapped cells meet unknown cells, and to reduce localization uncertainty, *place-revisiting actions* travel back to locations the robot has visited, or where it is able to observe a previously observed obstacle. The prevalence of these locations requires us to examine a large number of free grid cells to obtain a near-optimal solution.

Therefore, we define two sets of goal configurations in our action space. A set of frontier goal configurations $\mathcal{G}_{\text{frontier}}$ is used for exploration, and a set of place-revisiting goal configurations $\mathcal{G}_{\text{revisitation}}$ is used to correct the robot’s localization error. Algorithm 2 is designed for generating the frontiers, which are uniformly sampled on the boundary between the unknown and explored regions of the workspace. In Algorithm 3, the place-revisiting set of goal configurations is generated by choosing the candidates which can observe the largest number of occupied cells.

In Eq. (12), the log-determinant of the covariance matrix is derived from the M-step as our uncertainty metric. Since the estimated covariance has to be fused with a large initial

Algorithm 2 Frontier Goal Configurations

Require: Frontier cells \mathcal{F} , clearance map c , number of candidate goals N_f , minimum separation distance d

- 1: $\mathcal{G}_{\text{frontier}} \leftarrow \emptyset$
- 2: **while** $|\mathcal{G}_{\text{frontier}}| < N_f$ **do**
- 3: $\mathbf{f}^* \leftarrow \text{argmax}_{\mathbf{f}_i} c(\mathbf{f}_i)$
- 4: $\mathcal{G}_{\text{frontier}} \leftarrow \mathcal{G}_{\text{frontier}} \cup \{\mathbf{f}^*\}$
- 5: $\mathcal{F} \leftarrow \mathcal{F} \setminus \{\mathbf{f}_i \mid |\mathbf{f}_i - \mathbf{f}^*| \leq d\}$
- 6: **end while**
- 7: **return** Candidate exploration goals $\mathcal{G}_{\text{frontier}}$

Algorithm 3 Place-revisiting Goal Configurations

Require: Occupied cells \mathcal{O} , clearance map c , number of candidate goals N_r , revisitation radius r , minimum separation distance d

- 1: $\mathcal{G}_{\text{revisitation}} \leftarrow \emptyset$
- 2: $\mathcal{M} = \{\mathbf{m}_{1,\dots,k}\} \leftarrow \text{ComputeKMeans}(\mathcal{O}, k)$
- 3: **while** $|\mathcal{G}_{\text{revisitation}}| < N_r$ **do**
- 4: $\mathbf{m}^* = (x^*, y^*) \leftarrow \text{argmax}_{\mathbf{m}_i} |\mathcal{O}_i|$
- 5: $\mathcal{R}^* \leftarrow \{(x^* + r \cos(\theta), y^* + r \sin(\theta)) \mid \theta \in [0, 2\pi)\}$
- 6: $\mathbf{r}^* \leftarrow \text{argmax}_{\mathbf{r}_i \in \mathcal{R}^*} c(\mathbf{r}_i)$
- 7: $d^* \leftarrow \min_{\mathbf{g}_i \in \mathcal{G}_{\text{revisitation}}} |\mathbf{r}^* - \mathbf{g}_i|$
- 8: **if** $d^* \geq d$ **then**
- 9: $\mathcal{G}_{\text{revisitation}} \leftarrow \mathcal{G}_{\text{revisitation}} \cup \{\mathbf{r}^*\}$
- 10: **end if**
- 11: $\mathcal{M} \leftarrow \mathcal{M} \setminus \{\mathbf{m}^*\}$
- 12: **end while**
- 13: **return** Candidate revisitation goals $\mathcal{G}_{\text{revisitation}}$

covariance, the log-determinant, or D-optimality criterion, is guaranteed to be monotonically non-increasing during the exploration process, which is consistent with the conclusion in [22]. In addition to uncertainty criteria, it is valuable to incorporate a cost-to-go term to establish a trade-off between traveling cost and uncertainty reduction [5]. Thus, our utility function for motion planning is finalized as,

$$U_{\text{EM}}(X_{T:T+N}) = -\log \det(\tilde{\Sigma}_{X_{T:T+N}}) - \sum_k \log \det(\hat{\Sigma}_{\mathbf{v}_k}) - \alpha d(X_{T:T+N}), \quad (28)$$

where α is a weight on path distance $d(X_{T:T+N})$. In our experiments, we adopt a linearly decaying weight function with respect to traveled distance, whose parameters are determined experimentally and applied consistently throughout our algorithm comparisons below.

IV. EXPERIMENTAL RESULTS

Here we present experimental results from (1) a simulation of a sonar-equipped underwater robot exploring and mapping two planar, previously unmapped environments, and (2) the real autonomous exploration of a harbor environment with a BlueROV2 underwater robot. One of the simulation environments is populated with point landmarks, permitting a solution that employs landmark nodes in the SLAM factor

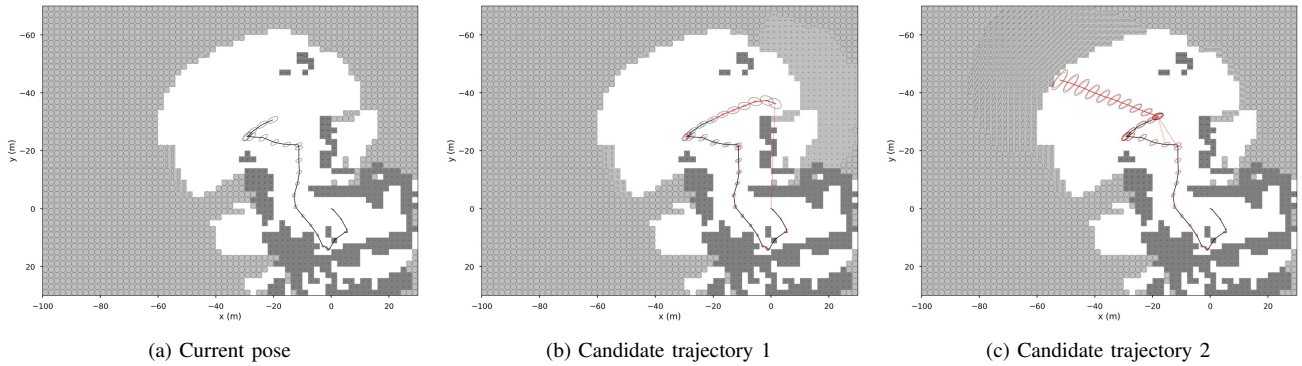


Fig. 2: Planning over virtual maps, with two candidate trajectories shown. Virtual landmarks are maintained within the cells of a low-resolution occupancy map (the higher-resolution occupancy maps used for motion planning and collision avoidance are shown in our video attachment). The error covariances of all occupied and unknown cells are shown as ellipses drawn inside each respective map cell, along with anticipated loop closure constraints. The maps shown were produced using real sonar data gathered by our ROV (Fig. 6).

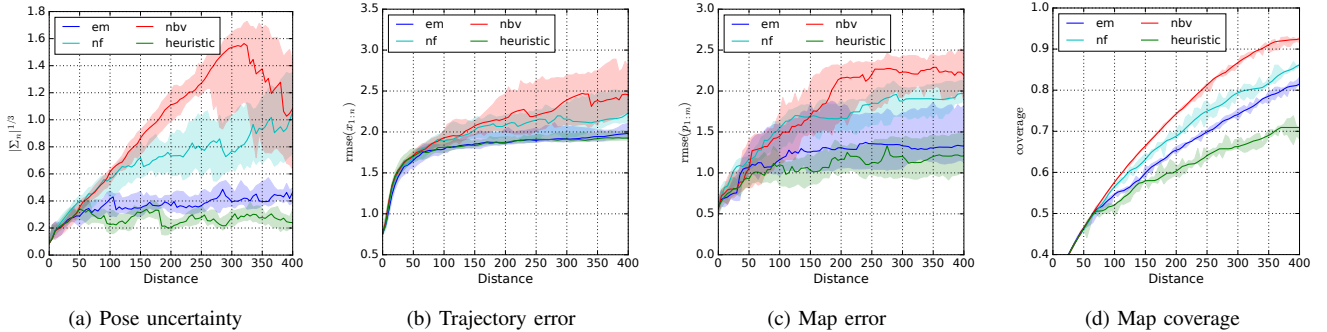


Fig. 3: Exploration performance mean values in the landmark environment of Fig. 4 (a) for 45 trials (five for each robot start location depicted in Fig. 4), with 95 percent confidence intervals shown. Values along the x-axis denote the robot’s distance traveled.

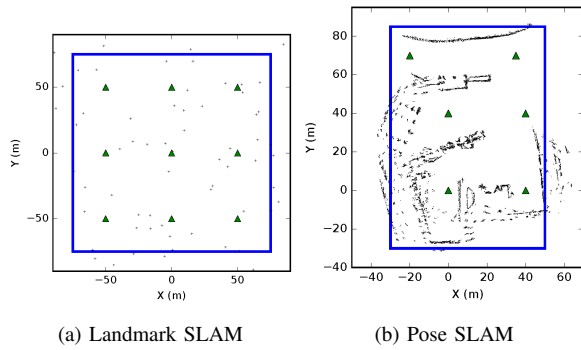


Fig. 4: Simulated exploration environments for landmark and pose SLAM. The exploration bounding box is denoted by a blue rectangle, and exploration is initialized from the green triangles. The robot may only travel within the bounding box, and only observations of the area within the bounding box are used to evaluate map coverage.

graph. The other environment is populated with heterogeneous structures, to which pose SLAM is applied instead.

The simulation environments were designed to emulate our subsequent real experiment, and are aimed at evaluating our algorithms quantitatively over a larger number of trials. The simulated sonar operates at 5 Hz and has a field of view with range $r = [0\text{ m}, 30\text{ m}]$ and horizontal aperture $\theta = [-65^\circ, 65^\circ]$. Since we assume a 2D environment, the vertical aperture is ignored in the simulation. Zero-mean Gaussian noise is added to range and bearing measurements: $\sigma_r = 0.2\text{ m}$, $\sigma_\theta = 0.02\text{ rad}$. We provide simulated odometry

measurements at 5 Hz, which can be obtained from DVL and IMU sensors in our real experiments. Additionally, we add zero-mean Gaussian noise to the 2D transform: $\sigma_x = \sigma_y = 0.08\text{ m}$, $\sigma_\theta = 0.003\text{ rad}$.

The two maps used for simulated experiments are shown in Figure 4. In the landmark map (containing a randomly generated collection of landmarks that is kept constant throughout all trials), the robot is deployed from nine different starting positions. In the pose SLAM map (whose structures are derived from our real-world experimental results), the robot is deployed from one of six starting positions. The exploration process is terminated when there are no frontiers remaining in the map that can be feasibly reached by our planner.

A. Algorithm Comparison

In the comparisons to follow, four algorithms are considered, all of which employ different techniques to repeatedly select one of the goal configurations generated by our motion planner described in Sec. III-G. First, we examine frontier-based exploration [14], which explores by repeatedly driving to the *nearest frontier* goal configuration. Secondly, we examine a next-best-view exploration framework that selects the goal configuration anticipated to achieve the largest *information gain* with respect to the robot’s occupancy map, in accordance with the technique proposed in [8]. Thirdly, we examine the active SLAM framework of [13]. When a robot’s *pose uncertainty* (per the D-optimality criterion [17])




Fig. 5: Exploration performance mean values in the simulated marina environment of Fig. 4 (b) for 60 trials (ten for each robot start location depicted in Fig. 4), with 95 percent confidence intervals shown. Values along the x-axis denote the robot’s distance traveled.

exceeds a designated threshold, this framework selects the place-revisiting goal configuration that maximizes a utility function expressing a weighted tradeoff between uncertainty reduction and information gain. At all other times, the approach reverts to that of [8], purely seeking information gain. In the results to follow, we will refer to this approach as the *heuristic* approach, due to its use of a tuned threshold on pose uncertainty to determine when place revisiting occurs. These three algorithms are compared against our proposed EM exploration algorithm. The heuristic and EM algorithms have been parameterized to achieve the closest possible equivalence in trajectory estimation error (the root mean square error of the full pose history, recorded at each instant in time during exploration), providing a baseline from which to examine other performance characteristics that differ.

B. Simulated Exploration over Landmarks

In the landmark environment shown in Figure 4a, we compared the nearest frontier (NF), next-best-view (NBV), heuristic, and EM algorithms across 45 simulated trials (five for each robot start location). In this environment, the landmarks are assumed to be infinitesimally small, and not to pose a collision hazard. The average performance of the four exploration algorithms is shown in Figure 3. The NBV method achieves superior coverage of the environment, but it also yields the highest pose uncertainty, trajectory error, and map error. The NF approach achieves poorer coverage than NBV, but achieves lower pose uncertainty and map/trajectory error during exploration. The heuristic approach covers the environment the least efficiently, while it offers the lowest pose uncertainty and map error during the exploration process. Our proposed EM algorithm achieves map coverage superior to the heuristic approach, with comparable but slightly higher pose uncertainty and map error.

C. Simulated Exploration with Pose SLAM

As there is no ground truth information available in our real-world experiments, we have instead created the map of Fig. 4b using manually collected data from the harbor environment used in those experiments. Specifically, the map of Fig. 4b is represented by an experimentally derived point cloud. During simulated exploration, feature points that fall within the sensor field of view are sampled and Gaussian noise is added - these sensor observations are used




Fig. 6: Our custom-instrumented BlueROV2 robot.

to populate an occupancy map. Ten trials of each algorithm were run at each of six different start locations (denoted as green triangles in Fig. 4). Mapping and localization errors are reported with respect to our experimentally-derived ground truth point cloud map. The mapping error is computed as follows: for every estimated feature point, we compute its distance to the nearest ground truth point.

The average performance of each algorithm is shown in Figure 5. As before, the NBV algorithm achieves the most efficient exploration, but at the expense of high pose uncertainty, map error, and trajectory error. The pose uncertainty, trajectory error, and map error of the heuristic and EM algorithms are similar in value, and significantly lower than that of the NF and NBV algorithms. While the heuristic and EM algorithms have similar pose uncertainty and map/trajectory error, the heuristic approach falls behind EM in terms of map coverage.

D. Exploration with a Real Underwater Robot

The hardware used to test this work is a modified BlueROV2, pictured in Fig. 6. This platform accommodates an Oculus M750d multi-beam imaging sonar, a Rowe SeaPilot DVL, a VectorNav VN100 IMU, and a Bar30 pressure sensor. Similar to our simulation, the sonar field of view is $r = [0 \text{ m}, 30 \text{ m}]$ and $\theta = [-65^\circ, 65^\circ]$, with a 20° vertical aperture. Although the pitch angle of the sonar is adjustable, in the experiments to follow, it is set to 0° . The sonar is operated at 750 MHz with 512 beams, 4 mm range resolution and 1° angular resolution. The ROV is operated at a fixed

depth of 1m throughout our experiments. Data is streamed via a tether to a surface laptop and processed in real-time on the laptop. The surface laptop is equipped with an Intel i7-4710 quad-core CPU, NVIDIA Quadro K1100M GPU, and 8GB of RAM. All components of the software stack are handled by the surface laptop, including image processing, the SLAM solution, belief propagation, motion planning, and issuing control inputs to the vehicle.

Real world experiments were carried out in a marina at the United States Merchant Marine Academy (USMMA). The goal was to explore and map the portions of the environment that lie within a bounding box of dimensions 130m \times 60m, similar to that depicted in Fig. 4b. Three runs were performed for each of the four competing algorithms, and all twelve trials can be viewed side-by-side in our video attachment.

V. CONCLUSION AND FUTURE WORK

In this paper we have presented and evaluated the concept of using *virtual maps*, comprised of uniformly discretized virtual landmarks, to support the efficient autonomous exploration of unknown environments while managing state estimation and map uncertainty. Our proposed EM exploration algorithm, when used in conjunction with a sonar keyframe-based pose SLAM architecture, equips an underwater robot with the ability to autonomously explore a cluttered underwater environment while maintaining an accurate map, pose estimate, and trajectory estimate, in addition to achieving a time-efficient coverage rate. The advantages of the proposed framework are demonstrated in simulation, and a qualitative proof of concept is presented over several real-world underwater exploration experiments. These experiments represent, to our knowledge, the first instance of fully autonomous exploration of a real-world, obstacle-filled outdoor environment in which an underwater robot directly incorporates its SLAM process and predictions based on that process into its decision-making. Future work entails the expansion of our framework to support three-dimensional sonar-based occupancy mapping.

ACKNOWLEDGMENTS

We would like to thank Arnaud Croux, Timothy Osedach and Stephane Vannuffelen of Schlumberger-Doll Research Center for guidance and support that was instrumental to the success of our field experiments. We also thank Prof. Carolyn Hunter of the U.S. Merchant Marine Academy for facilitating access to the marina used in the experiments.

REFERENCES

- [1] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- [2] J. Wang and B. Englot, "Autonomous exploration with expectation-maximization," in *Proceedings of the International Symposium on Robotics Research (ISRR)*, 2017.
- [3] J. Wang, T. Shan, and B. Englot, "Virtual maps for autonomous exploration with pose SLAM," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4899–4906.
- [4] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information based adaptive robotic exploration," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 1, 2002, pp. 540–545.
- [5] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using Rao-Blackwellized particle filters," in *Proceedings of Robotics: Science and Systems (RSS)*, 2005, pp. 65–72.
- [6] R. Valencia, J. Valls Miró, G. Dissanayake, and J. Andrade-Cetto, "Active pose SLAM," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 1885–1891.
- [7] B. Charrow, S. Liu, V. Kumar, and N. Michael, "Information-theoretic mapping using Cauchy-Schwarz quadratic mutual information," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4791–4798.
- [8] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon "next-best-view" planner for 3D exploration," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1462–1468.
- [9] C. Papachristos, S. Khattak, and K. Alexis, "Uncertainty-aware receding horizon exploration and mapping using aerial robots," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 4568–4575.
- [10] E. Vidal, J. D. Hernández, K. Istenič, and M. Carreras, "Online view planning for inspecting unexplored underwater structures," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1436–1443, 2017.
- [11] N. Palomeras, N. Hurtós, E. Vidal, and M. Carreras, "Autonomous exploration of complex underwater environments using a probabilistic next-best-view planner," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1619–1625, 2019.
- [12] S. M. Chaves, A. Kim, and R. M. Eustice, "Opportunistic sampling-based planning for active visual SLAM," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2014, pp. 3073–3080.
- [13] S. Suresh, P. Sodhi, J. Mangelson, D. Wetttergreen, and M. Kaess, "Active SLAM using 3D submap saliency for underwater volumetric exploration," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [14] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA)*, July 1997, pp. 146–151.
- [15] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "iSAM2: Incremental smoothing and mapping using the Bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [16] F. Dellaert and M. Kaess, "Square root SAM: Simultaneous localization and mapping via square root information smoothing," *The International Journal of Robotics Research*, vol. 25, no. 12, pp. 1181–1203, 2006.
- [17] H. Carrillo, I. Reid, and J. A. Castellanos, "On the comparison of uncertainty criteria for active SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2012, pp. 2080–2087.
- [18] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Robotics and Autonomous Systems*, vol. 57, no. 12, pp. 1198–1210, 2009.
- [19] V. Ila, L. Polok, M. Solony, P. Smrz, and P. Zemeik, "Fast covariance recovery in incremental nonlinear least square solvers," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4636–4643.
- [20] H. Li, F. Nashashibi, and M. Yang, "Split covariance intersection filter: Theory and its application to vehicle localization," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 4, pp. 1860–1871, 2013.
- [21] P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [22] H. Carrillo, Y. Latif, M. L. Rodriguez-Arevalo, J. Neira, and J. A. Castellanos, "On the monotonicity of optimality criteria during exploration in active SLAM," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 1476–1483.