Active SLAM over Continuous Trajectory and Control: A Covariance-Feedback Approach

Shumon Koga Arash Asgharivaskasi Nikolay Atanasov

***Abstract*— This paper proposes a novel active Simultaneous Localization and Mapping (SLAM) method with continuous trajectory optimization over a stochastic robot dynamics model. The problem is formalized as a stochastic optimal control over the continuous robot kinematic model to minimize a cost function that involves the covariance matrix of the landmark states. We tackle the problem by separately obtaining an open-loop control sequence subject to deterministic dynamics by iterative Covariance Regulation (iCR) and a closed-loop feedback control under stochastic robot and covariance dynamics by Linear Quadratic Regulator (LQR). The proposed optimization method captures the coupling between localization and mapping in predicting uncertainty evolution and synthesizes highly informative sensing trajectories. We demonstrate its performance in active landmark-based SLAM using relative-position measurements with a limited field of view.**

I. Introduction

Simultaneous Localization and Mapping (SLAM) has been instrumental for enabling autonomous robots to transition from controlled, structured, and fully known environments to operation in a priori unknown real-world conditions [7], [24]. Many current SLAM techniques, however, remain *passive* in their utilization of sensor data. Active SLAM [11] is an extension of the SLAM problem which couples perception and control, aiming to acquire more information about the environment and reduce the uncertainty in the localization and mapping process. Active SLAM introduces unique challenges related to keeping the map and location estimation processes accurate, and yet computing and propagating uncertainty over many potential sensing trajectories efficiently to select an informative one.

Many existing works in active perception decouple the localization and mapping problems and assume known robot states when planning active mapping trajectories. The literature on active mapping can be categorized according to the map representation it employs. Some techniques use volumetric mapping, which represents occupancy (e.g., occupancy grid) or obstacle distance (e.g., signed distance field) at a finite number of voxels obtained by discretizing the environment. Other techniques employ landmark-based mapping, which represents positions of a finite number of landmarks (e.g., objects or visual features) in the environment. While a volumetric representation captures the complete geometric structure of the environment, landmark-based mapping requires much less memory. One of the earliest approaches for active mapping [28] is based on detecting and planning a shortest path to frontiers (boundaries between explored and unexplored space) in a volumetric map. Information-theoretic planning is an alternative approach, which utilizes an information measure to quantify and minimize the uncertainty in the map, as developed first in [12] and subsequently used widely in robotics [20], [13], [14]. Efficient computation methods for computing uncertainty of volumetric maps have been proposed in [10] for Cauchy-Schwarz quadratic mutual information (CSQMI), and in [29] for fast Shannon mutual information (FSMI). Active mapping for truncated signed distance field (TSDF) reconstruction has been considered in [25] and multi-category semantic maps have been studied in [3]. Existing methods are, however, limited to discrete control spaces, typically with a finite number of possible control inputs, such as [25], [10], [29]. Recently, in [16], we have developed a continuous trajectory optimization method for active mapping, named iterative Covariance Regulation (iCR). We have introduced a differentiable field of view in sensing model, and apply gradient descent method for obtaining an open-loop control sequence to maximize the differential entropy of the map.

In applications other than mapping, motion planning under uncertainty in the robot states has been developed for several robotics tasks, such as reaching a goal without collisions with obstacles. To cope with the uncertainty in the motion model, the probability density function of the robot state given a sensory data is constructed, named as *belief space*. An important work in this area is Belief Roadmaps [22], which projects a roadmap from the state space to the belief space and seeks an optimal path in the constructed graph. Such a sampling-based method has been further developed in [2], [1]. Alternatively, continuous-space optimization methods have also been proposed for belief space planning. Representative work by [27] applies an iterative LQG [26], which computes both a nominal open-loop trajectory and a feedback control policy through iterative solutions of a dynamic programming. The authors have utilized iLQG for the belief dynamics consisting of the mean and the covariance of the robot state through EKF estimate utilizing a sensory data. However, the method in [27] limits the observation model to be smooth for enabling the gradient computation in iLQG, while a typical measurement by a camera field of view does not follow such a smooth observation model. Recently, [23] has relaxed the assumption by introducing a probabilistic visibility model in sensing, and proposed a novel motion planning method with guaranteeing a constraint on the uncertainty in the robot state to be satisfied, via

We gratefully acknowledge support from ARL DCIST CRA W911NF-17-2-0181 and NSF FRR CAREER 2045945.

The authors are with the Department of Electrical and Computer Engineering, UC San Diego, 9500 Gilman Drive, La Jolla, CA, 92093-0411, {skoga,aasghari,natanasov}@ucsd.edu.

employing an augmented Lagrangian method. This paper differs from the belief space planning proposed in [27], [23] in the sense that we take into account a target dynamics as a belief state and propose a method of solving an open-loop trajectory and closed-loop policy *separately*, which is computationally efficient since the iteration is needed only for an open-loop trajectory under a deterministic dynamics.

Active SLAM is a challenging problem due to the mutual dependence among the accuracy of the robot localization, the performance of the mapping, and the trade-off between exploring new areas and exploiting uncertainty minimization in visited areas. Several techniques in the literature approach the active SLAM problem with a landmark-based mapping using greedy planning over a discrete control space, see [8] for instance. To avoid the costly global planning with long time horizons, in [18], the authors have introduced an attractor in landmark-based active SLAM, which incorporates global information about the environment for the local planner, and applied a model predictive control approach. Utilizing the idea of introducing the attractor, in [5], a decentralized nonmyopic approach to a multi-robot landmark-based active SLAM has been proposed via exploiting the sparsity in the information filter. In [15], a multi-robot landmark-based active SLAM has been tackled by develoing a scalable sampling-based planning. While almost all of the literature in active SLAM have employed a discretized control space planning, [19] has developed an online path planning method for active SLAM under a continuous control space by a Bayesian optimization for updating a parameter in a control policy. However, the method requires a sampling of the state propagation to approximate a mean square error of SLAM and an iteration of policy search, which renders a difficulty in the real-time implementation with an online planning.

Unlike existing active SLAM techniques which commonly consider a discretized control space and known robot poses, this work develops an active SLAM method with continuous trajectory optimization over a stochastic robot dynamics model utilizing offline planning. A major advantage of the proposed method is its ability to capture the coupling between localization and mapping in predicting uncertainty evolution and to synthesize highly informative sensing trajectories due to the continuous-space optimization, which is computationally efficient and capable of real-time implementation through an offline computation. We first provide a general formulation of the active information acquisition as studied in [4], [25], [16]. Apart from the previous literature, we include a stochastic process noise in the robot dynamics, which needs to be dealt with in application to active SLAM. Next, we propose a method for obtaining a nominal open-loop trajectory via iCR for deterministic robot dynamics and a closed-loop control policy by LQR for a linearized stochastic system around the iCR trajectory. Then, we apply these techniques to active SLAM to estimate the positions of a finite number of landmarks.

II. Problem Statement

We consider a sensing system with state $\mathbf{x}_k \in \mathbb{R}^{n_x}$ and control input $\mathbf{u}_k \in \mathbb{R}^m$ at time $t_k \in \mathbb{R}_+$ where $\{t_k\}_{k=0}^K$ for some $K \in \mathbb{N}$ is an increasing sequence. The task is to reduce uncertainty in a target state, denoted as $\mathbf{y}_k \in \mathbb{R}^{n_y}$, through on-board sensors which receives a sensor measurement $\mathbf{z}_k \in \mathbb{R}^{n_z}$ as a function of both the sensing system and target states. We consider a stochastic nonlinear dynamics of the system state, a linear stochastic dynamics of the target state, and a linear observation model with respect to the target state, described by

$$\begin{aligned} \mathbf{x}_{k+1} &= \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k), \\ \mathbf{y}_{k+1} &= A\mathbf{y}_k + \Xi_k^{1/2}\boldsymbol{\xi}_k, \\ \mathbf{z}_k &= H(\mathbf{x}_k)\mathbf{y}_k + V(\mathbf{x}_k)^{1/2}\mathbf{v}_k, \end{aligned} \tag{1}$$

where $\mathbf{w}_k \sim \mathcal{N}(0, W_k)$, $\boldsymbol{\xi}_k \sim \mathcal{N}(0, I_{n_y})$, $\mathbf{v}_k \sim \mathcal{N}(0, I_{n_z})$, $W_k \in \mathbb{S}_{\succ 0}^{n_x \times n_x}$, $\Xi_k \in \mathbb{S}_{\succ 0}^{n_y \times n_y}$, $V : \mathbb{R}^{n_x} \to \mathbb{S}_{\succ 0}^{n_z \times n_z}$, and $H : \mathbb{R}^{n_x} \to \mathbb{R}^{n_z \times n_y}$. $\mathbb{S}_{\succ 0}^{n \times n}$ is a set of positive definite matrices within $\mathbb{R}^{n \times n}$.

Active information acquisition is a motion planning problem for a system dynamics aiming to minimize some uncertainty measure of the target state. One representative candidate for the measure is the differential entropy in the target state conditioned on the sensor states and measurements. In [17], for deterministic system dynamics (i.e., $\mathbf{w}_k = 0$), the problem of minimizing the differential entropy is shown to be equivalent to minimizing the log determinant of the covariance matrix of the target state as a deterministic optimal control, which is known as D-optimality in optimal experimental design [21]. There are several other criteria in optimal experimental design, such as A-optimality, minimizing the trace of the covariance, and E-optimality, minimizing the maximum eigenvalue [9]. Here, we consider a general cost function over the covariance matrix to capture all possible optimality criteria.

We approach the problem of minimizing an uncertainty criterion over the target state subject to motion and sensor models in (1) by extending the formulation in [4] to include stochastic noise in the robot dynamics. Let $\Sigma_k \in \mathbb{S}_{\succ 0}^{n_y \times n_y}$ be the covariance matrix of the target state. Then, the optimization problem in this paper is written as:

$$\min_{\mathbf{U} \in \mathbb{R}^{mK}} \mathbb{E}\left\{ \sum_{k=0}^{K-1} \bar{c}_k\left(\Sigma_k\right) + \bar{c}_K\left(\Sigma_K\right) \right\}, \tag{2}$$

$$\text{s.t.} \quad \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k),$$

$$\Sigma_{k+1} = A\left(\Sigma_k^{-1} + M(\mathbf{x}_{k+1})\right)^{-1} A^\top + \Xi_k, \tag{3}$$

$$M(\mathbf{x}) = H(\mathbf{x})^\top V(\mathbf{x})^{-1} H(\mathbf{x}), \tag{4}$$

where $\bar{c}_k : \mathbb{S}_{\succ 0}^{n_y \times n_y} \to \mathbb{R}$ for $k \in \{0, \ldots, K\}$ is an objective function with respect to the covariance matrix, $M : \mathbb{R}^{n_x} \to \mathbb{R}^{n_y \times n_y}$ is the so-called "sensor information matrix" . In D-optimal design to minimize the differential entropy, we can set $\bar{c}_k = 0$ for $k \in \{0, \ldots, K-1\}$ and $\bar{c}_K(\Sigma_K) = \log\det(\Sigma_K)$. In this paper, we keep setting a general objective function $\bar{c}_k$.

We tackle the optimal control problem by a separate planning of an open-loop trajectory optimization (iCR) developed in [16], and a closed-loop feedback control by LQR. While iCR in [16] has been developed for a robot motion model with an $SE(3)$ pose state represented by a matrix, a typical LQR approach is applicable to dynamics described with respect to a vector state. Therefore, we reformulate the Riccati update (3) of the covariance matrix by the dynamics of the vector state. Let $\boldsymbol{\sigma}_k \in \mathbb{R}^{n_\sigma}$, where $n_\sigma := n_y(n_y+1)/2$, be a vector representation of the covariance matrix Σ_k defined by $\boldsymbol{\sigma}_k := \mathrm{vech}(\Sigma_k)$, where $\mathrm{vech}(\cdot) : \mathbb{S}^{n_y \times n_y}_{\succ 0} \to \mathbb{R}^{n_y(n_y+1)/2}$ is a half-vectorization operator applied to a symmetric matrix. By reformulating the Riccati update in (3) and rewriting the cost function (2) with respect to the vector state $\boldsymbol{\sigma}$, the optimization problem we consider can be recast as follows.

Problem. Obtain a control policy $\mathbf{u}_k = \boldsymbol{\pi}_k(\mathbf{x}_k, \boldsymbol{\sigma}_k)$, where $\boldsymbol{\pi}_k : \mathbb{R}^{n_x+n_\sigma} \to \mathbb{R}^m$ for $k = 0, \ldots, K-1$ to solve the following stochastic optimal control problem:

$$\min_{\boldsymbol{\pi}_0,\ldots,\boldsymbol{\pi}_{K-1}} \mathbb{E}\left\{ \sum_{k=0}^{K-1} c_k(\boldsymbol{\sigma}_k) + c_K(\boldsymbol{\sigma}_K) \right\} \tag{5}$$

where $c_k : \mathbb{R}^{n_\sigma} \to \mathbb{R}$, subject to

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \boldsymbol{\pi}_k(\mathbf{x}_k, \boldsymbol{\sigma}_k), \mathbf{w}_k), \tag{6}$$

$$\boldsymbol{\sigma}_{k+1} = \mathbf{g}(\boldsymbol{\sigma}_k, \mathbf{x}_{k+1}). \tag{7}$$

III. Planning Method

This section proposes a method to solve (5)–(7) by designing an open-loop control sequence and closed-loop control policy separately.

A. Linearized system around iCR trajectory

iCR proposed by [16] provides the solution to the deterministic case (i.e., $\mathbf{w}_k = 0$) of the optimal control stated above, for active exploration and mapping with the cost of log determinant of the covariance matrix. Extending iCR to minimizing a general cost, we obtain a nominal open-loop trajectory $\{\bar{\mathbf{x}}_{k+1}, \bar{\mathbf{u}}_k, \bar{\boldsymbol{\sigma}}_{k+1}\}_{k=0}^{K-1}$ which is a solution to

$$[\bar{\mathbf{u}}_1, \ldots, \bar{\mathbf{u}}_{K-1}] = \underset{\bar{\mathbf{u}}_1,\ldots,\bar{\mathbf{u}}_{K-1}}{\mathrm{argmin}} \left(\sum_{k=0}^{K-1} c_k(\bar{\boldsymbol{\sigma}}_k) + c_K(\bar{\boldsymbol{\sigma}}_K) \right),$$
$$\text{s.t.} \quad \bar{\mathbf{x}}_{k+1} = \mathbf{f}(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0), \quad \bar{\boldsymbol{\sigma}}_{k+1} = \mathbf{g}(\bar{\boldsymbol{\sigma}}_k, \bar{\mathbf{x}}_{k+1}).$$

The solution is obtained via gradient-descent for a multi-step control sequence $\mathbf{U} = [\bar{\mathbf{u}}_1^\top, \ldots, \bar{\mathbf{u}}_{K-1}^\top]^\top$, via iterative update of the control sequence by $\mathbf{U} \leftarrow \mathbf{U} - \boldsymbol{\alpha} \frac{\partial J^{\mathbf{U}}}{\partial \mathbf{U}}$ with a step size $\boldsymbol{\alpha}$ and the cost $J^{\mathbf{U}} = \sum_{k=0}^{K-1} c_k(\boldsymbol{\sigma}_k) + c_K(\boldsymbol{\sigma}_K)$, where the gradient $\frac{\partial J^{\mathbf{U}}}{\partial \mathbf{U}}$ is computed analytically. Around the nominal open-loop trajectory $\{\bar{\mathbf{x}}_{k+1}, \bar{\mathbf{u}}_k, \bar{\boldsymbol{\sigma}}_{k+1}\}_{k=0}^{K-1}$ and the mean of the noise $\mathbf{w}_k = 0$, the nonlinear stochastic dynamics (6), (7) can be described by a linear time-varying system as a first-order approximation through Taylor expansion (see [30]). Let us define the error variables: $\tilde{\mathbf{x}}_k = \mathbf{x}_k - \bar{\mathbf{x}}_k, \tilde{\boldsymbol{\sigma}}_k = \boldsymbol{\sigma}_k - \bar{\boldsymbol{\sigma}}_k, \tilde{\boldsymbol{\pi}}_k(\tilde{\mathbf{x}}_k, \tilde{\boldsymbol{\sigma}}_k) = \boldsymbol{\pi}_k(\mathbf{x}_k, \boldsymbol{\sigma}_k) - \bar{\mathbf{u}}_k$. Then, linearizing the dynamics (6), (7) around the nominal trajectory $\{\bar{\mathbf{x}}_{k+1}, \bar{\boldsymbol{\sigma}}_{k+1}\}_{k=0}^{K-1}$, the dynamics for the error variables are

$$\tilde{\mathbf{x}}_{k+1} = E_k \tilde{\mathbf{x}}_k + B_k \tilde{\boldsymbol{\pi}}_k(\tilde{\mathbf{x}}_k, \tilde{\boldsymbol{\sigma}}_k) + D_k \mathbf{w}_k,$$
$$\tilde{\boldsymbol{\sigma}}_{k+1} = F_k \tilde{\boldsymbol{\sigma}}_k + G_k(E_k \tilde{\mathbf{x}}_k + B_k \tilde{\mathbf{u}}_k + D_k \mathbf{w}_k),$$

where

$$E_k = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0)}, B_k = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right|_{(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0)}, D_k = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{w}}\right|_{(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0)}, \tag{8}$$

$$F_k = \left.\frac{\partial \mathbf{g}}{\partial \boldsymbol{\sigma}}\right|_{(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0)}, \quad G_k = \left.\frac{\partial \mathbf{g}}{\partial \mathbf{x}}\right|_{(\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k, 0)}. \tag{9}$$

We aim to minimize the cost (5) subject to the linearized stochastic dynamics by designing a control policy $\tilde{\boldsymbol{\pi}}_k(\mathbf{x}, \boldsymbol{\sigma})$. Applying the Taylor expansion to (5) around the nominal trajectory and approximating by second-order accuracy yields

$$c_k(\boldsymbol{\sigma}_k) \approx \tilde{\boldsymbol{\sigma}}_k^\top C_k \tilde{\boldsymbol{\sigma}}_k + \bar{\mathbf{b}}_k^\top \tilde{\boldsymbol{\sigma}}_k + c_k(\bar{\boldsymbol{\sigma}}_k), \tag{10}$$

$$C_k := \frac{\partial^2 c_k}{\partial \boldsymbol{\sigma}^2}(\bar{\boldsymbol{\sigma}}_k), \quad \bar{\mathbf{b}}_k^\top := \frac{\partial c_k}{\partial \boldsymbol{\sigma}}(\bar{\boldsymbol{\sigma}}_k). \tag{11}$$

At the same time, to validate the linearization, the trajectories should stay around the nominal trajectories, namely, the error variables should stay around zero. Pursuing these two objectives, and introducing the new state variable $\mathbf{s}_k \in \mathbb{R}^{n_x+n_\sigma}$ defined by $\mathbf{s}_k = [\tilde{\mathbf{x}}_k^\top, \tilde{\boldsymbol{\sigma}}_k^\top]^\top, \tilde{\boldsymbol{\pi}}_k(\mathbf{s}_k) = \boldsymbol{\pi}_k(\mathbf{x}_k, \boldsymbol{\sigma}_k) - \bar{\mathbf{u}}_k$ the dynamics and the cost function can be written with respect to $\mathbf{s}_k$ as

$$\mathbf{s}_{k+1} = \mathcal{A}_k \mathbf{s}_k + \mathcal{B}_k \tilde{\boldsymbol{\pi}}_k(\mathbf{s}_k) + \mathcal{D}_k \mathbf{w}_k, \tag{12}$$

$$\tilde{J} = \mathbb{E}\left\{ \sum_{k=0}^{K-1} (\mathbf{s}_k^\top \mathcal{Q}_k \mathbf{s}_k + \mathbf{b}_k^\top \mathbf{s}_k + \tilde{\boldsymbol{\pi}}_k(\mathbf{s}_k) R_k \tilde{\boldsymbol{\pi}}_k(\mathbf{s}_k)) \right\}$$
$$+ \mathbb{E}\left\{ \mathbf{s}_K^\top \mathcal{Q}_K \mathbf{s}_K + \mathbf{b}_K^\top \mathbf{s}_K \right\}, \tag{13}$$

where $n_s = n_x + n_\sigma$, $\mathcal{A}_k \in \mathbb{R}^{n_s \times n_s}$, $\mathcal{B}_k \in \mathbb{R}^{n_s \times m}$, $\mathcal{D}_k \in \mathbb{R}^{n_s \times n_x}$, $\mathcal{Q} \in \mathbb{R}^{n_s \times n_s}$, $\mathbf{b} \in \mathbb{R}^{n_s}$ are defined by

$$\mathcal{A}_k = \begin{bmatrix} E_k & 0 \\ G_k E_k & F_k \end{bmatrix}, \mathcal{B}_k = \begin{bmatrix} B_k \\ G_k B_k \end{bmatrix}, \mathcal{D}_k = \begin{bmatrix} D_k \\ G_k D_k \end{bmatrix}, \tag{14}$$

$$\mathcal{Q}_k = \begin{bmatrix} Q_k^{(1)} & 0 \\ 0 & Q_k^{(2)} + C_k \end{bmatrix}, \quad \mathbf{b}_k = \begin{bmatrix} \mathbf{0} \\ \bar{\mathbf{b}}_k \end{bmatrix}, \tag{15}$$

where $Q_k^{(1)} \in \mathbb{R}^{n_x \times n_x}, Q_k^{(2)} \in \mathbb{R}^{n_\sigma \times n_\sigma}$, and $R_k \in \mathbb{R}^{m \times m}$ are weight matrices to be determined by the user.

B. LQR-based closed-loop feedback control

We derive a closed-loop feedback control policy to minimize (13) subject to (12). Note that, unlike the standard LQR, the cost function (13) includes a linear term $\mathbf{b}_k^\top \mathbf{s}_k$, in both stage and terminal costs. Even then, one can show that the optimal control policy is a linear feedback control with a time-varying constant term, as stated below.

Proposition 1. *The closed-loop control*

$$\tilde{\boldsymbol{\pi}}_k(\mathbf{s}_k) = L_k^* \mathbf{s}_k + \varepsilon_k^*, \tag{16}$$

$$L_k^* = -(R + \mathcal{B}_k^\top P_{k+1} \mathcal{B}_k)^{-1} \mathcal{B}_k^\top P_{k+1} \mathcal{A}_k, \tag{17}$$

$$\varepsilon_k^* = -\frac{1}{2}(R + \mathcal{B}_k^\top P_{k+1} \mathcal{B}_k)^{-1} \mathcal{B}_k^\top \mathbf{d}_{k+1}, \tag{18}$$

where $P_k \in \mathbb{R}^{n_s \times n_s}$, $\mathbf{d}_k \in \mathbb{R}^{n_s}$, and $\delta_k \in \mathbb{R}$ at $k = K$ are given by $P_K = \mathcal{Q}_K, \mathbf{d}_K = \mathbf{b}_K, \delta_K = 0$, and recursively updated from $k+1$ to k for $k = K-1, K-2, \ldots, 0$ as follows:

$$P_k = \mathcal{Q}_k + \mathcal{A}_k^\top P_{k+1} \mathcal{A}_k - \mathcal{A}_k^\top P_{k+1} \mathcal{B}_k (R + \mathcal{B}_k^\top P_{k+1} \mathcal{B}_k)^{-1} \mathcal{B}_k^\top P_{k+1} \mathcal{A}_k, \quad (19)$$

$$\mathbf{d}_k = \mathbf{b}_k + \mathcal{A}_k^\top \mathbf{d}_{k+1} - \mathcal{A}_k^\top P_{k+1} \mathcal{B}_k (R + \mathcal{B}_k^\top P_{k+1} \mathcal{B}_k)^{-1} \mathcal{B}_k^\top \mathbf{d}_{k+1}, \quad (20)$$

$$\delta_k = \delta_{k+1} + \mathrm{tr}(\mathcal{D}_k^\top P_{k+1} \mathcal{D}_k W_k) - \frac{1}{4} \mathbf{d}_{k+1}^\top \mathcal{B}_k (R + \mathcal{B}_k^\top P_{k+1} \mathcal{B}_k)^{-1} \mathcal{B}_k^\top \mathbf{d}_{k+1} \quad (21)$$

minimizes the cost function (13) *subject to the system dynamics* (12) *with optimal cost:*

$$\min_{\tilde{\mathbf{u}}_0, \ldots, \tilde{\mathbf{u}}_{K-1}} \tilde{J} = V_0(\mathbf{s}) = \mathbf{s}^\top P_0 \mathbf{s} + \mathbf{d}_0^\top \mathbf{s} + \delta_0,$$

for a given initial condition $\mathbf{s}_0 = \mathbf{s}$.

The proof is done using a dynamic programming method, and is omitted in this paper due to space constraints.

IV. Application to Active SLAM

We apply the proposed planning method to an active SLAM problem, where a set of landmarks in the environment is regarded as the target state. The task is to estimate the landmark positions and the pose of a sensing robot, and to plan motion that reduces the uncertainty in these estimates.

A. Differential-drive motion model and its linearization

Let $\mathbf{x} = [\mathbf{p}^\top, \theta]^\top \in \mathbb{R}^3$ be the state of a ground robot, where $\mathbf{p} \in \mathbb{R}^2$ is the robot position and $\theta \in [-\pi, \pi)$ is the robot's heading angle. Let $\mathbf{u} \in [v, \omega] \in \mathbb{R}^2$ be the robot's control input, where v is the linear velocity and ω is the angular velocity. We model the robot dynamics $\mathbf{f} : \mathbb{R}^3 \times \mathbb{R}^2 \times \mathbb{R}^3 \to \mathbb{R}^3$ in (6) using a differential-drive kinematic model with time discretization τ:

$$\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k) = \mathbf{x}_k + \tau \begin{bmatrix} v_k \mathrm{sinc}\,(a_k) \cos(\theta_k + a_k) \\ v_k \mathrm{sinc}\,(a_k) \sin(\theta_k + a_k) \\ \omega_k \end{bmatrix} + \mathbf{w}_k, \quad (22)$$

where $a_k := \frac{\omega_k \tau}{2}$.

B. Limited field-of-view sensing model

The landmarks are modeled as points $\mathbf{y}^{(j)} \in \mathbb{R}^2$ in the environment, for all $j \in \{1, \ldots, n_l\}$, where n_l is a number of landmark. Noting that the landmark location is static, the matrices in the target dynamics given in (1) are set as:

$$A = I_{2n_l \times 2n_l}, \quad \Xi_k = 0_{2n_l \times 2n_l}. \quad (23)$$

For a given robot state $\mathbf{x} \in \mathbb{R}^3$ and the position $\mathbf{y}^{(j)} \in \mathbb{R}^2$ of the j-th landmark, we consider the robot body-frame coordinates of $\mathbf{y}^{(j)}$:

$$\mathbf{q}\left(\mathbf{x}, \mathbf{y}^{(j)}\right) = R^\top(\theta)(\mathbf{y}^{(j)} - \mathbf{p}), \quad (24)$$

where $R : \mathbb{R} \to SO(2) \subset \mathbb{R}^{2\times 2}$ is a 2-D rotation matrix of the robot pose. Here, we define the set of indices of the landmarks within the field of view $\mathcal{F} \subset \mathbb{R}^2$ of the robot, as follows:

$$\mathcal{I}_{k,\mathcal{F}} = \left\{ \forall j \in \{1, \ldots, n_l\} \middle| \mathbf{q}\left(\mathbf{x}_k, \mathbf{y}^{(j)}\right) \in \mathcal{F} \right\} \quad (25)$$

We suppose to have both range and bearing measurements, which capture the relative landmark positions in robot frame as follows:

$$\mathbf{z}_k = \left[\ \{\bar{\mathbf{z}}(\mathbf{x}_k, \mathbf{y}^{(j)})\}_{j \in \mathcal{I}_{k,\mathcal{F}}} \ \right] \in \mathbb{R}^{2|\mathcal{I}_{k,\mathcal{F}}|}, \quad (26)$$

$$\bar{\mathbf{z}}(\mathbf{x}, \mathbf{y}^{(j)}) = \mathbf{q}(\mathbf{x}, \mathbf{y}^{(j)}) + \Gamma^{1/2}(\mathbf{x}, \mathbf{y}^{(j)}) \mathbf{v}, \quad (27)$$

where $\Gamma : \mathbb{R}^3 \times \mathbb{R}^2 \to \mathbb{S}^{2\times 2}_{\succ 0}$ is the sensor noise covariance and $\mathbf{v} \sim \mathcal{N}(0, I_2)$. Notice that the sensing model (27) is nonlinear with respect to the target state, while up to the previous sections we have considered a linear sensing model. In the application to active SLAM in this section, we deal with the nonlinear sensing model by employing the linearization around some estimate of the target state, and utilize it for both planning and SLAM.

C. Differentiable field of view

One special characteristic of the sensing model (26) with a limited field of view (FoV) is that the measurement dimension is dependent on the robot pose state at time k. This is caused by a binary (observable or unobservable) sensing within the FoV, which makes the sensing model non-differentiable with respect to the state and hence the control input. To deal with the challenge, we use a differentiable FoV proposed in [16], where the measurement is supposed to be obtained for all landmark states, while the noise covariance is supposed to become approximately infinity outside the FoV. Namely, the measurement function $\mathbf{h} : \mathbb{R}^3 \times \mathbb{R}^{2n_l} \to \mathbb{R}^{2n_l}$ and the noise covariance matrix are formalized as:

$$\mathbf{h}(\mathbf{x}, \mathbf{y}) = [\mathbf{q}(\mathbf{x}, \mathbf{y}^{(1)})^\top, \ldots, \mathbf{q}(\mathbf{x}, \mathbf{y}^{(n_l)})^\top]^\top, \quad (28)$$

$$V(\mathbf{x}) = \mathrm{diag}(\bar{V}(\mathbf{x}, \hat{\mathbf{y}}^{(1)}), \ldots, \bar{V}(\mathbf{x}, \hat{\mathbf{y}}^{(n_l)})) \in \mathbb{R}^{2n_l \times 2n_l},$$

where $\hat{\mathbf{y}}^{(j)} \in \mathbb{R}^2$ is an initial estimate of the j-th landmark state, $\bar{V}(\mathbf{x}, \mathbf{y}^{(j)}) : \mathbb{R}^3 \times \mathbb{R}^2 \to \mathbb{R}^{2\times 2}$ is a noise covariance of j-th landmark with differentiable FoV, given by

$$\bar{V}(\mathbf{x}, \mathbf{y}^{(j)}) = \left(1 - \Phi(d(\mathbf{q}(\mathbf{x}, \mathbf{y}^{(j)}), \mathcal{F}))\right)^{-1} \Gamma(\mathbf{x}, \mathbf{y}^{(j)}),$$

where $\Phi : \mathbb{R} \to [0,\ 1]$ is the Gaussian CDF defined by $\Phi(x) = \frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{x}{\sqrt{2}\kappa} - 2\right)\right]$ and $d(\mathbf{q}, \mathcal{F})$ is a signed distance function defined below.

Definition 1. The *signed distance function* $d : \mathbb{R}^2 \to \mathbb{R}$ associated with a set $\mathcal{F} \subset \mathbb{R}^2$ is:

$$d(\mathbf{q}, \mathcal{F}) = \begin{cases} -\min_{\mathbf{q}^* \in \partial\mathcal{F}} \|\mathbf{q} - \mathbf{q}^*\|, & \text{if} \quad \mathbf{q} \in \mathcal{F}, \\ \min_{\mathbf{q}^* \in \partial\mathcal{F}} \|\mathbf{q} - \mathbf{q}^*\|, & \text{if} \quad \mathbf{q} \notin \mathcal{F}, \end{cases} \quad (29)$$

where $\partial\mathcal{F}$ is the boundary of $\mathcal{F}$.

A 2-D plot of $1 - \Phi(d(\mathbf{q}(\mathbf{x}, \mathbf{y}^{(j)}), \mathcal{F})$, which is the amplification factor in the differentiable FoV, is shown in

[16]. In accordance with the linear sensing model in (1), we take a linearization of (28). Let $H : \mathbb{R}^3 \to \mathbb{R}^{2n_l \times 2n_l}$ be the measurement matrix defined by $H(\mathbf{x}) = \frac{\partial \mathbf{h}}{\partial \mathbf{y}}(\mathbf{x}, \hat{\mathbf{y}})$. Then, with (28), one can show H is a block diagonal matrix, given by $H(\mathbf{x}) = \mathrm{diag}(\bar{H}(\mathbf{x}_k, \hat{\mathbf{y}}^{(1)}), \ldots, \bar{H}(\mathbf{x}_k, \hat{\mathbf{y}}^{(n_l)}))$, where $\bar{H}(\mathbf{x}_k, \mathbf{y}^{(j)}) : \mathbb{R}^3 \times \mathbb{R}^2 \to \mathbb{R}^{2\times 2}$ is the measurement matrix for j-th landmark at time k, given as $\bar{H}(\mathbf{x}_k, \mathbf{y}^{(j)}) := \frac{\partial \mathbf{q}}{\partial \mathbf{y}^{(j)}}(\mathbf{x}_k, \mathbf{y}^{(j)}) = R^\top(\theta_k)$, where we utilized (24). Since the product of the block diagonal matrix with same size is also a block diagonal matrix of each product, the sensor information matrix defined by (4) is also a block diagonal matrix, given by $M(\mathbf{x}) = \mathrm{diag}(\bar{M}(\mathbf{x}, \hat{\mathbf{y}}^{(1)}), \ldots, \bar{M}(\mathbf{x}, \hat{\mathbf{y}}^{(n_l)}))$, where $\bar{M}(\mathbf{x}, \mathbf{y}^{(j)}) = \left(1 - \Phi(d(\mathbf{q}^{(j)}, \mathcal{F}))\right) R(\theta)\Gamma^{-1}R^\top(\theta)$. Since the matrix M and the initial covariance matrix are block diagonal matrices with same size, and the fact that the landmark is static (23), Riccati update (3) leads to that the covariance matrix Σ_k is also a block diagonal matrix with the same size for all k, i.e., $\Sigma_k = \mathrm{diag}(\Sigma_k^{(1)}, \Sigma_k^{(2)}, \cdots, \Sigma_k^{(n_l)})$, satisfying

$$\Sigma_{k+1}^{(j)} = \left(\left(\Sigma_k^{(j)}\right)^{-1} + \bar{M}\left(\mathbf{x}_{k+1}, \hat{\mathbf{y}}_0^{(j)}\right)\right)^{-1} \tag{30}$$

D. iCR for active SLAM with trace minimization

Following [16], in the derivation of iCR, the robot pose dynamics is described by $SE(2)$ pose kinematics, in which the robot state is defined in $SE(2) \subset \mathbb{R}^{3\times 3}$. The $SE(2)$ pose kinematics is equivalent to the differential-drive motion model in (22), through the exponential map $T = \exp(\hat{\mathbf{x}})$ and the logarithm map $\mathbf{x} = \log(T)^\vee$ [6]. While in [16] the cost is set to a log determinant of the covariance matrix at final time, in this paper we consider the cost to be a trace of the covariance matrix, which eases the computation of the gradient (11) of the cost with respect to the vectorized covariance. Describing the robot dynamics by $SE(2)$ pose kinematics, and taking into account the differentiable field of view formulation presented above, iCR for active SLAM problem is developed, thereby the nominal trajectory $\{\bar{T}_{0:K}, \bar{\Sigma}_{0:K}, \bar{\mathbf{u}}_{0,K-1}\}$ for a deterministic $SE(2)$ pose kinematics is obtained.

E. LQR gain derivation

As presented in Secion III, for deriving the LQR gains, provided the nominal iCR trajectory, Jacobian matrices (8) (9) must be computed. As stated above, iCR trajectory $\{\bar{T}_{0:K}, \bar{\Sigma}_{0:K}, \bar{\mathbf{u}}_{0,K-1}\}$ is obtained for a $SE(2)$ pose state $\bar{T} \in SE(2) \subset \mathbb{R}^{3\times 3}$ and the covariance matrix $\bar{\Sigma} \in \mathbb{S}_{\succ 0}^{2n_l \times 2n_l}$, which need to be converted to vectorized states. The robot vector state $\mathbf{x} \in \mathbb{R}^3$ is obtained by $\mathbf{x} = \log(T)^\vee$. Regarding the covariance vector state $\boldsymbol{\sigma}$, in Section II, we define by the half-vectorization operator $\mathrm{vech}(\cdot)$ as a general case. However, in active SLAM we consider in this section, owing to the independency among covariances of each landmark as shown in (30) and its symmetric property, we can define a covariance vector state with a lower dimension than the half-vectorization operator as follows. Let $\boldsymbol{\sigma}^{(j)} \in \mathbb{R}^3$ be defined by $\boldsymbol{\sigma}^{(j)} = \left[\mathbf{e}_1^\top \Sigma^{(j)} \mathbf{e}_1, \mathbf{e}_1^\top \Sigma^{(j)} \mathbf{e}_2, \mathbf{e}_2^\top \Sigma^{(j)} \mathbf{e}_2\right]^\top$,

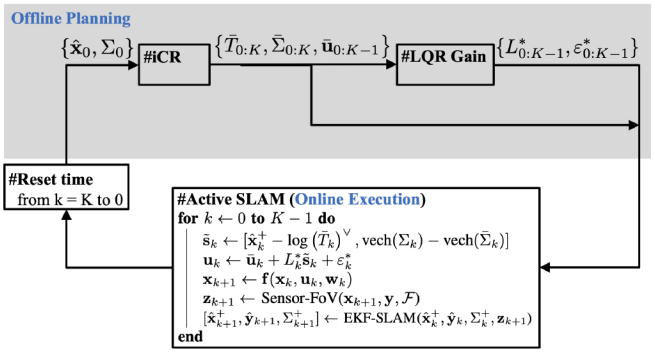

Fig. 1: Diagram of active SLAM via iCR-LQR.

where $\Sigma^{(j)} \in \mathbb{S}_{\succ 0}^{2\times 2}$ is a j-th block diagonal matrix in $\Sigma \in \mathbb{S}_{\succ 0}^{2n_l \times 2n_l}$. Then, we define the covariance vector state by $\boldsymbol{\sigma} := [\boldsymbol{\sigma}^{(1)}, \boldsymbol{\sigma}^{(2)}, \ldots, \boldsymbol{\sigma}^{(n_l)}] \in \mathbb{R}^{3n_l}$, which includes all the variables in the covariance matrix obeying the Riccati update. Hereafter, this conversion is denoted as $\boldsymbol{\sigma} = \mathrm{vecbl}(\Sigma)$.

Through this conversion from the covariance matrix to the covariance vector state, we can derive the Jacobian matrix analytically. The computation steps of obtaining the LQR gain is shown in Algorithm 1, which includes "Jacob-robot" as a Jacobian matrices of the robot dynamics (22), and "Jacob-Riccati" as a Jacobian matrices of the Riccati update provided in Appendix. "Grad-cost" computes the gradient of the cost function with respect to the vectorized covariance state as given in (11). Considering the cost of minimizing a trace of the covariance, the variables (11) in the cost function are obtained as $C_k = 0_{3n_l \times 3n_l}$ and $\bar{\mathbf{b}}_k = [\boldsymbol{\eta}^\top, \boldsymbol{\eta}^\top, \ldots, \boldsymbol{\eta}^\top]^\top \in \mathbb{R}^{3n_l}$ where $\boldsymbol{\eta} = [1, 0, 1]^\top \in \mathbb{R}^3$, since each $\boldsymbol{\eta}$ in $\bar{\mathbf{b}}_k$ corresponds to the diagonal element of $\Sigma^{(j)}$ with respect to the vectorized covariance $\boldsymbol{\sigma}^{(j)}$.

Algorithm 1: LQR gain

Data: iCR trajectory $\bar{T}_{1:K}, \bar{\mathbf{u}}_{1:K-1}, \bar{\Sigma}_{1:K}$, weight matrices $Q_k^{(1)}, Q_k^{(2)}, R_k$
$C, \bar{\mathbf{b}} \leftarrow$ Grad-cost($\bar{\boldsymbol{\sigma}}_K$) in (11)
$[P, \mathbf{d}, \delta] \leftarrow [\mathcal{Q}_K, \mathbf{b}_K, 0]$ from (15)
Solve LQR gain backward in time
for $k \leftarrow K-1$ **to** 0 **do**
  $[\bar{\mathbf{x}}_k, \bar{\mathbf{x}}_{k+1}, \bar{\boldsymbol{\sigma}}_k] \leftarrow$
  $[\log(\bar{T}_k)^\vee, \log(\bar{T}_{k+1})^\vee, \mathrm{vecbl}(\bar{\Sigma}_k)]$
  $E, B, D \leftarrow$ Jacob-robot($\bar{\mathbf{x}}_k, \bar{\mathbf{u}}_k$)
  $F, G \leftarrow$ Jacob-Riccati($\bar{\boldsymbol{\sigma}}_k, \bar{\mathbf{x}}_{k+1}$)
  $\mathcal{A}, \mathcal{B}, \mathcal{D} \leftarrow$ (14) using $[E, B, D, F, G]$
  $C, \bar{\mathbf{b}} \leftarrow$ Grad-cost($\bar{\boldsymbol{\sigma}}_k$) in (11)
  $\mathcal{Q}, \mathbf{b} \leftarrow$ (15) using $[Q_k^{(1)}, Q_k^{(2)}, C, \bar{\mathbf{b}}]$
  $[L_k^*, \varepsilon_k^*] \leftarrow$ (17), (18) using $[P, \mathbf{d}, \mathcal{A}, \mathcal{B}, R]$
  $\mathbf{d} \leftarrow$ (20) using $[\mathbf{d}, \mathbf{b}, P, \mathcal{A}, \mathcal{B}, R]$
  $\delta \leftarrow$ (21) using $[\delta, P, \mathbf{d}, \mathcal{B}, R, \mathcal{D}, W_k]$
  $P \leftarrow$ (19) using $[P, \mathcal{Q}, \mathcal{A}, \mathcal{B}, R]$
end
Output: $L_{0:K-1}^*, \varepsilon_{0:K-1}^*$

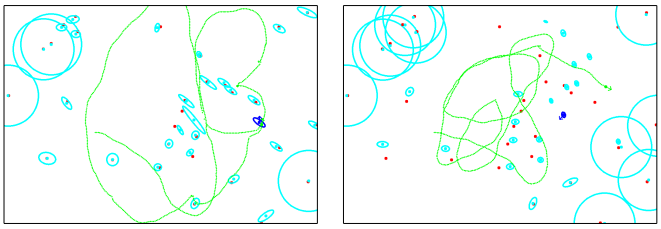

(a) Closed-loop iCR + LQR (b) Open-loop iCR

Fig. 2: Active SLAM via closed-loop and open-loop control policies. The green dotted line shows the robot ground truth trajectory. The blue dot shows robot pose, while the surrounding ellipse corresponds to the covariance of robot position. Red dots indicate the ground truth landmark positions. Cyan dots and ellipses indicate the mean and covariance of landmarks, respectively.

F. EKF-SLAM

We construct the estimator of both the landmark position and the robot position by Extended Kalman Filter (EKF). The probabilistic landmark position and the robot position are set as a Gaussian distribution: $\begin{bmatrix} \mathbf{x}_k \\ \mathbf{y} \end{bmatrix} | \mathbf{z}_{1:k} \sim \mathcal{N}\left(\begin{bmatrix} \hat{\mathbf{x}}_k^+ \\ \hat{\mathbf{y}}_k \end{bmatrix}, \Sigma_k^+, \right)$ for posteriori estimate, and $\begin{bmatrix} \mathbf{x}_{k+1} \\ \mathbf{y} \end{bmatrix} | \mathbf{z}_{1:k} \sim \mathcal{N}\left(\begin{bmatrix} \hat{\mathbf{x}}_{k+1}^- \\ \hat{\mathbf{y}}_k \end{bmatrix}, \Sigma_{k+1}^-, \right)$ for a priori estimate. Once the measurement is obtained as given in (26), which is a measured landmark relative position in robot-body frame within FoV, we reconstruct the measured state as $\mathbf{z}_k = [\mathbf{z}_k^{(1)}, \mathbf{z}_k^{(2)}, \ldots, \mathbf{z}_k^{(n_l)}] \in \mathbb{R}^{2n_l}$, where

$$\mathbf{z}_k^{(j)} = \begin{cases} \bar{\mathbf{z}}(\mathbf{x}_k, \mathbf{y}^{(j)}), & \text{if} \quad j \in \mathcal{I}_{k,\mathcal{F}} \\ \mathbf{q}(\hat{\mathbf{x}}_k^-, \hat{\mathbf{y}}_{k-1}^{(j)}), & \text{otherwise} \end{cases} \tag{31}$$

Using the reconstructed sensor state (31), we implement EKF for SLAM by updating the mean and covariance of both priori and posteriori estimates. Note that, since the innovation term in EKF is $\mathbf{z}_k - \mathbf{h}(\hat{\mathbf{x}}_k^-, \hat{\mathbf{y}}_{k-1})$, where $\mathbf{h}$ is given by (28), the reconstructed sensor state (31) makes the innovation term zero in j-th landmark estimate for all $j \notin \mathcal{I}_{k,\mathcal{F}}$. Namely, all the landmark estimate outside FoV does not have an update through applying (31) to EKF-SLAM.

A diagram depicting the structure of the entire proposed algorithm is shown in Fig. 1.

G. Evaluation

We examine the performance of the proposed method in a simulated environment with dimensions 100[m] $\times$ 70[m], where the landmarks are located following a uniform random distribution. The robot follows the $SE(2)$ motion model of (22) with $W_k = \mathrm{diag}\,(0.1, 0.1, 0.01)$, and its on-board sensor measures the relative position of visible landmarks in the robot frame. The field of view $\mathcal{F}$ is set as an isosceles triangle with height 20 [m] and the angle between the two legs equal to 120°. The relative position measurements are corrupted by an additive Gaussian noise with zero mean and covariance $\Gamma = \mathrm{diag}\,(0.1, 0.1)$. The control $\mathbf{u}_k$ and measurements $\mathbf{z}_{k+1}$ are given to EKF-SLAM for state estimation, where we assume the noise covariances W_k and Γ are known to EKF-SLAM. We initialize the mean $\left[\hat{\mathbf{x}}_0^+ \hat{\mathbf{y}}_0^+\right]^\top$ with the ground truth position of the robot and landmarks, added by a Gaussian noise with variance 25 [m^2], while the state covariance is initialized as $\Sigma_0^+ = 25 I_{n_s \times n_s}$. During each planning phase, we begin by computing the initial iCR control sequence $\bar{\mathbf{u}}_{0:K-1}$ for planning horizon $K = 5$, where the differentiable FoV is parameterized by $\kappa = 10$ and the gradient decent update is done for 10 iterations with $\boldsymbol{\alpha} = \mathrm{diag}\,(0.005, 0.0005)$. The obtained control sequence $\bar{\mathbf{u}}_{0:K-1}$ is regularized by LQR where $Q_k^{(1)} = \mathrm{diag}\,(10, 10, 1)$, $Q_k^{(2)} = I_{n_l \times n_l} \otimes \mathrm{diag}\,(1, 0.1, 1)$, and $R_k = \begin{bmatrix} 20 & 5 \\ 5 & 10 \end{bmatrix}$. The closed-loop control $\mathbf{u}_k$ is applied to the robot for K steps while the state mean and covariance are updated on each step. Fig. 2 shows examples of active SLAM using open-loop and closed-loop control policies. We observe that the LQR closed-loop control allows for larger exploration of the environment since the trajectory is constantly corrected by LQR, while for the case of open-loop policy, landmark entropy increases after execution of the control sequence, which encourages re-visiting the nearby landmarks and resulting in limited exploration.

Fig. 3 summarizes simulation results for a random policy, open-loop control obtained from iCR, and LQR closed-loop control over iCR output. Both open-loop and closed-loop policies outperform the random policy; however, the policy regulated by LQR shows more long-term stability and uncertainty reduction. This can be directly attributed to the cost function for LQR, where stability and landmark position uncertainty are explicitly factored in the model.

V. Conclusion

This paper developed a method for continuous trajectory optimization for active information acquisition problems. The problem is formalized as a stochastic optimal control problem to minimize an uncertainty measure over the target state. The novelty of the proposed method lies in (i) taking into account the process noise in robot dynamics, (ii) introducing a differentiable field of view for enabling the gradient computation, and (iii) planning an open-loop trajectory by iCR and a closed-loop control policy by LQR applied to a linearized system around iCR trajectory. We demonstrated the efficacy of the proposed method in a simulation of landmark-based active SLAM, aiming to map the landmarks and localize robot accurately.

References

[1] Ali-akbar Agha-mohammadi, Saurav Agarwal, Sung-Kyun Kim, Suman Chakravorty, and Nancy M Amato. SLAP: Simultaneous localization and planning under uncertainty via dynamic replanning in belief space. *IEEE Transactions on Robotics*, 34(5):1195–1214, 2018.

[2] Ali-Akbar Agha-Mohammadi, Suman Chakravorty, and Nancy M Amato. FIRM: Feedback controller-based information-state roadmap – a framework for motion planning under uncertainty. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4284–4291, 2011.

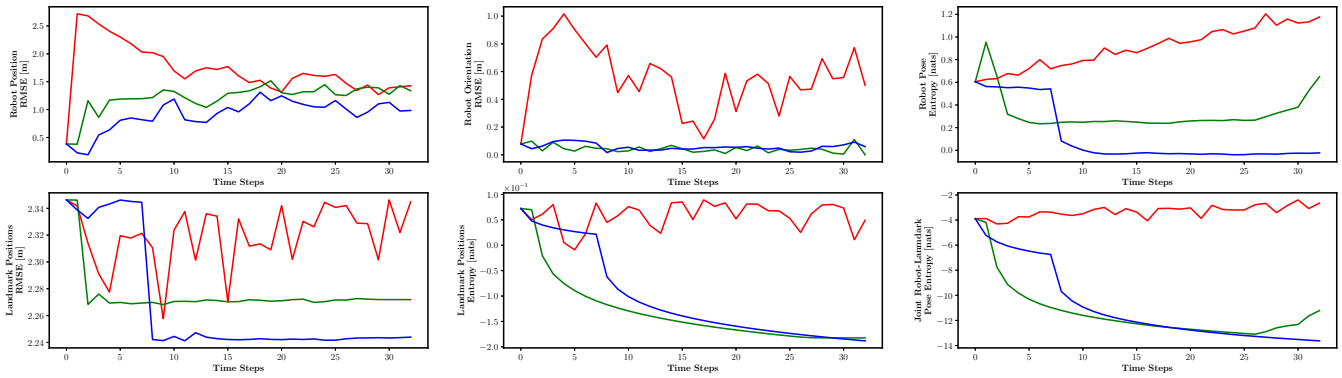

Fig. 3: Simulation results for random policy (red), iCR (green), and iCR with closed-loop LQR control (blue). Top three figures illustrate the root mean square error (RMSE) of robot position and orientation, as well as robot pose entropy. The bottom figures show RMSE of landmark positions, average landmark position entropy, and the joint robot-landmark pose entropy. The results are averaged over 5 random environments in which all three methods were executed.

[3] Arash Asgharivaskasi and Nikolay Atanasov. Active Bayesian multi-class mapping from range and semantic segmentation observation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.

[4] Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Information acquisition with sensing robots: Algorithms and error bounds. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 6447–6454, 2014.

[5] Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Decentralized active information acquisition: Theory and application to multi-robot SLAM. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4775–4782, 2015.

[6] Timothy D Barfoot. *State estimation for robotics*. Cambridge University Press, 2017.

[7] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Transactions on Robotics*, 32(6):1309–1332, 2016.

[8] Luca Carlone, Jingjing Du, Miguel Kaouk Ng, Basilio Bona, and Marina Indri. Active SLAM and exploration with particle filters using Kullback-Leibler divergence. *Journal of Intelligent & Robotic Systems*, 75(2):291–311, 2014.

[9] Henry Carrillo, Ian Reid, and José A Castellanos. On the comparison of uncertainty criteria for active SLAM. In *IEEE International Conference on Robotics and Automation*, pages 2080–2087, 2012.

[10] Benjamin Charrow, Sikang Liu, Vijay Kumar, and Nathan Michael. Information-theoretic mapping using cauchy-schwarz quadratic mutual information. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4791–4798, 2015.

[11] Jeffrey Delmerico, Elias Mueggler, Julia Nitsch, and Davide Scaramuzza. Active autonomous aerial exploration for ground robot path planning. *IEEE Robotics and Automation Letters*, 2(2):664–671, 2017.

[12] Alberto Elfes. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In *International Workshop on Reasoning with Uncertainty in Robotics*, pages 91–130. Springer, 1995.

[13] Ben Grocholsky. Information-theoretic control of multiple sensor platforms. *Ph.D. dissertation, Univ. of Sydney, 2002 [Online]. Available: http://www.acfr.usyd.edu.au*, 2002.

[14] Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.

[15] Y. Kantaros, B. Schlotfeldt, N. Atanasov, and G. J. Pappas. Asymptotically optimal planning for non-myopic multi-robot information gathering. In *Robotics: Science and Systems (RSS)*, 2019.

[16] Shumon Koga, Arash Asgharivaskasi, and Nikolay Atanasov. Active exploration and mapping via iterative covariance regulation over continuous $SE(3)$ trajectories. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.

[17] Jerome Le Ny and George J Pappas. On trajectory optimization for active sensing in Gaussian process models. In *Proceedings of the 48h IEEE Conference on Decision and Control (CDC) held jointly with 2009 28th Chinese Control Conference*, pages 6286–6292. IEEE, 2009.

[18] Cindy Leung, Shoudong Huang, and Gamini Dissanayake. Active SLAM using model predictive control and attractor based exploration. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5031. IEEE, 2006.

[19] Ruben Martinez-Cantin, Nando De Freitas, Eric Brochu, José Castellanos, and Arnaud Doucet. A Bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. *Autonomous Robots*, 27(2):93–103, 2009.

[20] Stewart J Moorehead, Reid Simmons, and William L Whittaker. Autonomous exploration using multiple sources of information. In *IEEE International Conference on Robotics and Automation (ICRA)*, volume 3, pages 3098–3103, 2001.

[21] A. Pázman. *Foundations of Optimum Experimental Design*. New York, NY, USA: Springer, 1986.

[22] Samuel Prentice and Nicholas Roy. The belief roadmap: Efficient planning in belief space by factoring the covariance. *The International Journal of Robotics Research*, 28(11-12):1448–1465, 2009.

[23] Shatil Rahman and Steven L Waslander. Uncertainty-constrained differential dynamic programming in belief space for vision based robots. *IEEE Robotics and Automation Letters*, 6(2):3112–3119, 2021.

[24] David M Rosen, Kevin J Doherty, Antonio Terán Espinoza, and John J Leonard. Advances in inference and representation for simultaneous localization and mapping. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:215–242, 2021.

[25] Kelsey Saulnier, Nikolay Atanasov, George J Pappas, and Vijay Kumar. Information theoretic active exploration in signed distance fields. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4080–4085, 2020.

[26] Emanuel Todorov and Weiwei Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *American Control Conference (ACC)*, pages 300–306, 2005.

[27] Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278, 2012.

[28] Brian Yamauchi. A frontier-based approach for autonomous exploration. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, pages 146–151, 1997.

[29] Zhengdong Zhang, Theia Henderson, Sertac Karaman, and Vivienne Sze. FSMI: Fast computation of shannon mutual information for information-theoretic mapping. *The International Journal of Robotics Research*, 39(9):1155–1177, 2020.

[30] Dongliang Zheng, Jack Ridderhof, Panagiotis Tsiotras, and Ali-akbar Agha-mohammadi. Belief space planning: A covariance steering approach. *arXiv preprint arXiv:2105.11092*, 2021.