Multi-Robot Autonomous Exploration and Mapping Under Localization Uncertainty with Expectation-Maximization

Yewei Huang[1] and Brendan Englot[1]

***Abstract*— We propose an autonomous exploration algorithm designed for distributed multi-robot teams, which takes into account map and localization uncertainties of range-sensing mobile robots. Virtual landmarks are used to quantify the combined impact of process noise and sensor noise on map uncertainty. Additionally, we employ an iterative expectation-maximization inspired algorithm to assess the potential outcomes of both a local robot's and its neighbors' next-step actions. The results of our experiments demonstrate the algorithm's capacity to strike a balance between curbing map uncertainty and achieving efficient task allocation among robots.**

I. Introduction

Autonomous exploration and mapping describes a single robot or a group of robots navigating themselves in an unknown or partially known environment without human intervention. An accurate environment map constructed by the robots serves as the fundamental basis for all subsequent specific robot tasks [1]. While the maturity of autonomous exploration for ground and aerial robots has been notably demonstrated, its application within marine environments remains a subject of ongoing inquiry [2]. In spite of the widespread coverage of Global Navigation Satellite System (GNSS) signals in most maritime regions, there exist many marine locations where GNSS signals are vulnerable due to obstruction or attenuation. In these areas, robots face a higher risk of collisions, emphasizing the critical need for an accurate environment map. However, deploying exploration in such contexts is a particular challenge for robot teams. This is primarily due to the high uncertainty introduced by unique environmental factors pertaining to the marine environment, especially when operating below the ocean's surface.

Autonomous exploration and mapping have been a vigorously discussed subject for several decades. Early research has primarily centered around optimizing task distribution among team members [3], [4]. As Simultaneous Localization and Mapping (SLAM) techniques have advanced, various approaches [5], [6], [7], [8], [9] have incorporated the concept of map uncertainty into autonomous exploration. These strategies consider Gaussian noise sensor models and Gaussian noise kinematic models, aiming to strike a balance between exploration efficiency and managing uncertainties in the resulting maps. Although these techniques may differ in terms of their map representation and merging techniques, they share a similar utility function that takes into account both the information gain of the latest generated map and the required travel distance to the waypoints under consideration. However, in the process of computing information gain, these techniques either ignore uncertainty propagation or have only propagated uncertainty locally, which is inadequate for situations with high localization uncertainty.

Building upon our previous work on single robot expectation-maximization (EM) inspired exploration [10], we introduce an asynchronous EM exploration algorithm for both centralized and distributed multi-robot teams. The algorithm is tightly coupled with a factor graph SLAM system. Upon a robot's arrival at the target position while the exploration remains ongoing, a virtual map is constructed based on the SLAM result to gauge the prevailing map uncertainty. Subsequently, for each potential new target, we execute an expectation-maximization procedure to assess the potential information gain associated with the forthcoming actions of the robot and its interactions with neighboring robots. Selection of a new target goal position is then determined by considering both the collective information gain of the entire team within the virtual map, and the efficiency of task allocation among robots.

Our research introduces several innovative contributions:

- An asynchronous multi-robot exploration framework catering to both centralized and distributed SLAM systems, taking into account efficient task allocation for exploration and addressing map uncertainty.
- Incorporating an expectation-maximization inspired technique to assess the future impact and interactions of a robot with its neighboring entities.
- Introducing an efficient inter-robot and local map uncertainty propagation approach, tailored to scenarios involving multiple robots and localization uncertainty.

II. Problem Formulation and Approach

We address an autonomous exploration problem that is tightly coupled with a SLAM factor graph for a team of n robots. We make the assumption that the initial states of all robots are sufficiently close to each other, enabling mutual observation among group members and facilitating an efficient map initialization process. Additionally, we impose a boundary on the exploration task, where the exploration process terminates upon fully exploring the enclosed environment.

A. Simultaneous Localization and Mapping

Let $N = \{1, 2, \cdots, n\}$ be the set of n robots. For each robot $\alpha \in N$, we denote its state at timestamp i as $\mathbf{x}_{\alpha,i}$. The robot odometry observation between present state $\mathbf{x}_{\alpha,i}$ and previous state $\mathbf{x}_{\alpha,i-1}$ is described by the equation:

$$\mathbf{z}_{\alpha,i}^{\alpha,i-1} = f(\mathbf{x}_{\alpha,i-1}, \mathbf{x}_{\alpha,i}) + \epsilon_{\alpha,i}^{\alpha,i-1}. \tag{1}$$

*This work was supported by ONR Grant N00014-21-1-2161.

[1]Y. Huang and B. Englot are with Stevens Institute of Technology, Hoboken, NJ, USA 07030. {yhuang85, benglot}@stevens.edu

Assume robot α observes a landmark state $\mathbf{l}_j$ at timestamp i, we can describe the landmark observation:

$$\mathbf{z}_j^{\alpha,i} = g(\mathbf{x}_{\alpha,i}, \mathbf{l}_j) + \epsilon_j^{\alpha,i}. \tag{2}$$

If another robot β is observed at timestamp i by robot α,we refer to this as the robot rendezvous observation:

$$\mathbf{z}_{\beta,j}^{\alpha,i} = f(\mathbf{x}_{\alpha,i}, \mathbf{x}_{\beta,j}) + \epsilon_{\beta,j}^{\alpha,i}, \tag{3}$$

where $f(\cdot)$ denotes the state transformation between robot states, $g(\cdot)$ is the state transformation from robot state to landmark state, and $\epsilon_{\alpha,i}^{\alpha,i-1}$, $\epsilon_j^{\alpha,i}$ and $\epsilon_{\beta,j}^{\alpha,i}$ are zero-mean Gaussian noise variables.

At present timestamp t, $\mathcal{X} = \{\mathbf{x}_i^\alpha | \alpha \in N, i \in [0,t]\}$ represents the set containing the states of all n robots from the initial timestamp 0 to the present timestamp t. $\mathcal{L} = \{\mathbf{l}_0, \mathbf{l}_1, \ldots, \mathbf{l}_m\}$ is the set of landmarks observed until timestamp t. Additionally, let $\mathcal{Z}$ be the set containing odometry observations, landmark observations and robot rendezvous observations from all timestamps. The SLAM problem can be represented as a maximum a posteriori estimation problem [11]:

$$\mathcal{X}^*, \mathcal{L}^* = \arg\max_{\mathcal{X},\mathcal{L}} P(\mathcal{X}, \mathcal{L}|\mathcal{Z}). \tag{4}$$

B. Expectation-Maximization Exploration

During the exploration process, when a robot α reaches its current target state $\mathbf{a}_\alpha^{\text{this}}$, it becomes necessary to determine a new target state $\mathbf{a}_\alpha^{\text{next}}$. Building upon our previous research on single-robot exploration [10], we consider a frontier-based strategy that incorporates two key factors: efficient task allocation among robots [3] and the maintenance of a relatively low uncertainty map.

To assess the map uncertainty, we create a virtual map $\mathcal{V}$ using the estimated robot and landmark states obtained from the SLAM step. For a comprehensive explanation of the virtual map construction process, please refer to our prior work [12]. When a target robot α reaches its target state $\mathbf{a}_\alpha^{\text{this}}$, we define:

$$\mathcal{X}^{\text{new}} = \mathcal{X}^{\text{old}} \cup \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}}. \tag{5}$$

With $\mathcal{X}^{\text{old}}$ containing the historical states of all n robots, $\mathcal{X}^{\text{predict}}$ denoting the set of predicted robot states for each current target state $\{\mathbf{a}_i^{\text{this}} | i \in N, i \neq \alpha\}$ and $\mathcal{X}^{\text{next}}$ representing the robot state sequence of robot α resulting from $\mathbf{a}_\alpha^{\text{next}}$. We first construct a virtual map using the historical data of robot teams:

$$\mathcal{V}^* = \arg\max_{\mathcal{V}} P(\mathcal{V}|\mathcal{X}^{\text{old}}, \mathcal{Z}^{\text{old}}). \tag{6}$$

Here, $\mathcal{Z}^{\text{old}}$ represents the observations associated with $\mathcal{X}^{\text{old}}$. Then we perform another maximum a posteriori estimation to estimate $\mathcal{X}^{\text{new}}$:

$$\mathcal{X}^{\text{new}*} = \arg\max_{\mathcal{X}^{\text{new}}} P(\mathcal{X}^{\text{new}}|\mathcal{V}^*, \mathcal{Z}^{\text{new}}), \tag{7}$$

$$\mathcal{Z}^{\text{new}} = \mathcal{Z}^{\text{old}} \cup \mathcal{Z}^{\text{predict}}. \tag{8}$$

The updated observation set, $\mathcal{Z}^{\text{new}}$, is a combination of both previously gathered observations, $\mathcal{Z}^{\text{old}}$, and anticipated future observations, $\mathcal{Z}^{\text{predict}}$. These predicted observations are determined via the measurement models outlined in Eq. (1), Eq. (2), and Eq. (3), based on the virtual map $\mathcal{V}^*$ and the target state assigned to each robot. Thus, we assess the overall local map uncertainty of robot α by computing the sum of the individual uncertainties for each map element in the virtual map $\mathcal{V}(\cdot)$ derived from Eq. (6) considering only the optimized states of robot α, $\mathcal{X}_\alpha^{\text{new}*}$:

$$U_M = \phi(\Sigma_{\mathcal{V}(\mathcal{X}_\alpha^{\text{new}*}, \mathcal{Z}^{\text{new}})}) \tag{9}$$

$$= \sum_{\mathbf{v_i} \in (\mathcal{X}_\alpha^{\text{new}*}, \mathcal{Z}^{\text{new}})} \phi(\Sigma_{\mathbf{v}_i}). \tag{10}$$

Each element $\mathbf{v}_i$ contributes to the sum based on its covariance Σ_{v_i}, and in this paper, we employ the D-Optimality metric as our uncertainty criterion ϕ.

To optimize the distribution of exploration tasks among the robots, we utilize a distance evaluation metric inspired by the methodology presented in [3]. For each potential new target state $\mathbf{a}_\alpha^{t\,\prime}$ assigned to robot α at time t, we quantify the effectiveness of task allocation using U_T:

$$U_T = \sum_{\mathbf{a}_\beta^i \in \mathcal{A}, \beta \neq \alpha} h(\|\mathbf{a}_\alpha^{t\,\prime} - \mathbf{a}_\beta^i\|_2), \tag{11}$$

$$h(d) = \begin{cases} 1 - \frac{d}{d_{\max}} & d < d_{\max} \\ 0 & d \geq d_{\max} \end{cases}. \tag{12}$$

Here, $\|\cdot\|_2$ represents the L2 norm, and $\mathcal{A}$ denotes the collection of all historical target states for all robots. Based on this, we can define the new target state for robot α as one of the potential target states that optimally balances the factors U_M, U_T and the Euclidean distance to target state factor, U_D, with scale factors λ_0, λ_1, λ_2:

$$\mathbf{a}_\alpha^{\text{next}} = \arg\min_{\mathbf{a}_\alpha^{t\,\prime}} (\lambda_0 U_M + \lambda_1 U_T + \lambda_2 U_D). \tag{13}$$

III. Algorithm and Results

In this section, we discuss the details of the proposed expectation-maximization algorithm. We consider a situation where a group of robots collaborates, and each individual robot within the group is furnished with both motion and perception sensors. A factor-graph based SLAM algorithm is operating at a specific frequency, concurrently generating a virtual map for each robot based on the trajectory from its SLAM solution. The robot team shares their optimized SLAM trajectories, landmark positions and history of chosen target states among its members. However, individual robots maintain virtual maps locally using the latest SLAM estimates.

After a robot α successfully reaches its current target state, its virtual map is updated and a group of potential new target states is chosen from the virtual map. Subsequently, the virtual observation is synthesized utilizing the prevailing environmental knowledge. Next, the graph-based uncertainty propagation computes the potential impact of the target state under consideration, leading to an update in the virtual map's covariances. The new target goal is then selected according to the utility function (Eq. (13)). Finally, the motion planner is initiated to formulate a series of actions leading robot α to the new target state.

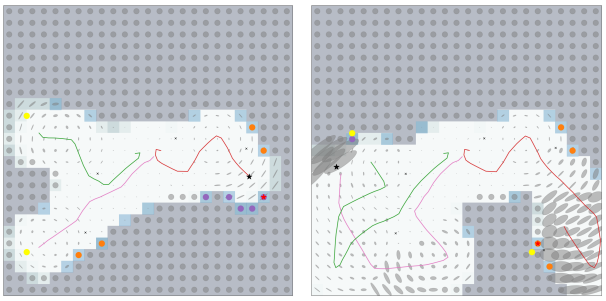

(a) An example of choosing an exploring frontier (purple) due to the current map uncertainty being relatively low.

(b) An example of choosing an revisiting frontier (orange) driven by the significant uncertainty in the map.

Fig. 1: In the $100m \times 100m$ virtual map created by the robot team, gray ellipses depict the uncertainty of each cell. The current robot position is denoted by a black star, and the newly selected target state is represented by a red star. Landmarks are expressed by black crosses. Potential frontiers emerge along the boundary between the explored and unexplored areas. The three types of frontiers—exploring, revisiting, and rendezvous—are differentiated by their respective colors: purple, orange, and yellow.

A. Virtual Map

As depicted in Eq. (6), the virtual map $\mathcal{V}$ of a finite environment is generated from the robot states $\mathcal{X}$ and their associated observations $\mathcal{Z}$. Assuming that $\mathcal{V}$ comprises b map cells $\mathbf{v}_i$ referred to as virtual landmarks, the posterior can be redefined as follows:

$$P(\mathcal{V}|\mathcal{X},\mathcal{Z}) = \prod_{\mathbf{v}_i \in \mathcal{V}} P(\mathbf{v}_i|\mathcal{X},\mathcal{Z}). \tag{14}$$

The likelihood of virtual landmark $\mathbf{v}_i$ being observed is $q(\mathbf{v}_i) = \mathbb{E}[P(\mathbf{v}_i|\mathcal{X},\mathcal{Z})]$. When this virtual landmark is observed by multiple robot states, the procedure for updating $q(\mathbf{v}_i)$ is the same as updating map cell values in an occupancy grid map [13]. We assume the virtual landmark $\mathbf{v}_i$ is observed by a robot state $\mathbf{x}_{\alpha,j}$, with its estimated value denoted as $\hat{\mathbf{x}}_{\alpha,j}$, and a marginal covariance of $\Sigma_{\mathbf{x}_{\alpha,j}}$. We can compute the covariance: $\Sigma_{\mathbf{v}_i} = \mathbf{H} \cdot \Sigma_{\mathbf{x}_{\alpha,j}} \cdot \mathbf{H}^\intercal$. Here, $\mathbf{H} = \frac{\partial g(\mathbf{x}_{\alpha,j},\mathbf{v}_i)}{\partial \mathbf{x}_{\alpha,j}}|\hat{\mathbf{x}}_{\alpha,j}$ represents the Jacobian matrix obtained by differentiating the landmark observation model $g(\mathbf{x}_{\alpha,j},\mathbf{v}_i)$ with respect to estimated robot state $\hat{\mathbf{x}}_{\alpha,j}$. We utilize Covariance Intersection to compute the covariance of a virtual landmark that is observed by multiple robot states, as detailed in [12]. Fig. 1 shows inter-robot virtual maps built from both local robot states and neighbors' robot states received by local robot α. The observed regions are highlighted in white; gray ellipses show covariances describing the uncertainty of the map's cells.

B. Uncertainty Propagation

As depicted in Fig. 1, most potential target states are chosen from the perimeters of the observed regions. Three types of frontiers are identified for selection: exploration frontiers close to the robot's latest position, revisiting frontiers near previously visited landmarks, and rendezvous frontiers, which are the current target positions of neighboring robots. Subsequently, for each potential target, a set of waypoints $\mathcal{X}^{\text{new}}$ is uniformly sampled along the shortest path connecting

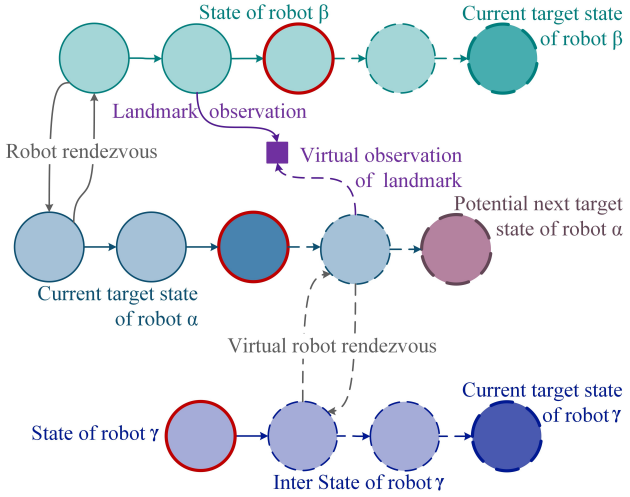

Fig. 2: Graph-based uncertainty propagation with virtual observations. Nodes representing the current robot states are distinguished by red edges. For every potential next target state (shown in pink), a trajectory simulation is executed, leading to the generation of a sequence of virtual observations indicated by dashed arrows. Simultaneously, we model the future states and observations of other robots as they approach their individual current target states.

Algorithm 1: Uncertainty Propagation

Global: Latest SLAM Graph $\mathcal{G}$ and Inter-robot Virtual Map $\mathcal{V}^*$
Input: Potential frontier state $\mathbf{a}_\alpha^{t\,\prime}$, Current robot states $\mathcal{X}^t$, Current Target States of Neighbors $\mathcal{A}^t$
Output: Optimized robot states $\mathcal{X}^{\text{new}*}$
$\mathcal{X}^{\text{predict}} \leftarrow \emptyset$
foreach $\mathbf{a}_\gamma^t \in \mathcal{A}^t, \mathbf{x}_\gamma^t \in \mathcal{X}^t$ **do**
  $\mathcal{X}_\gamma^{t:*} \leftarrow$ `GenerateVirtualWaypoints`$(\mathbf{a}_\gamma^t, \mathbf{x}_\gamma^t)$
  $\mathcal{X}^{\text{predict}} \leftarrow \mathbf{X}_\gamma^t \cup \mathcal{X}^{\text{predict}}$
end
$\mathcal{X}^{\text{next}} \leftarrow$ `GenerateVirtualWaypoints`$(\mathbf{a}_\alpha^{t\,\prime}, \mathbf{x}_\gamma^t)$
Calculate virtual observations
$\mathcal{Z}^{\text{predict}} \leftarrow$ `VirtualObserve` $(\mathcal{V}^*, \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}})$
Graph optimization
`UpdateGraph` $(\mathcal{G}, \mathcal{X}^{\text{predict}} \cup \mathcal{X}^{\text{next}})$
$\mathcal{X}^{\text{new}*} \leftarrow$ `OptimizeGraph`$(\mathcal{G})$
return $\mathcal{X}^{\text{new}*}$

the robot's current state and the next potential target state. Additionally, waypoints $\mathcal{X}^{\text{predict}}$ are generated to connect the present states of neighboring robots with their respective target states. This augmentation enhances the simulation of future steps. The process of generating virtual observations $\mathcal{Z}^{\text{predict}}$ along these waypoints is depicted in Fig. 2. Virtual odometry measurements are created between adjacent waypoints using Eq. (1). For virtual landmark observations, Eq. (2) is employed, generating observations between previously observed landmarks and nearby waypoints. When two robots are within each other's sensing range at the same timestep, a virtual robot observation is produced using Eq. (3). Following this, virtual observations are added into the SLAM graph to propagate uncertainty. The step-by-step procedure is outlined comprehensively in Algorithm 1.

Algorithm 2: Calculate U_M of potential frontier $\mathbf{a}_\alpha^{t\prime}$

Input: Inter-robot Virtual Map $\mathcal{V}^*$, Local Virtual Map $\mathcal{V}_\alpha^{\text{new}*}$
Output: Map uncertainty utility factor U_M
$U_M \leftarrow 0$
for $i \leftarrow 0$ **to** b **do**
  $\mathbf{v}_{i,g} \in \mathcal{V}^*$, $\mathbf{v}_{i,l} \in \mathcal{V}_\alpha^{\text{new}*}$
  # $q_{\min}$: minimum accepted probability of observed.
  if $q(\mathbf{v}_{i,l}) > q_{\min}$ ***and*** $q(\mathbf{v}_{i,g}) > q_{\min}$ **then**
    $U_M \leftarrow U_M + |\Sigma_{\mathbf{v}_{i,l}}|$
  end
end
return U_M

C. Map Uncertainty Utility Factor

We compute the uncertainty of map cells by considering both the likelihood of the existing virtual map $\mathcal{V}^*$ and the optimized robot states of robot α, $\mathcal{X}_\alpha^{\text{new}*} \subset \mathcal{X}^{\text{new}*}$. Referencing Fig. 3(f), it is evident that despite the pink robot's trajectory being impacted by localization uncertainty due to accumulated odometry error, this uncertainty is partially mitigated by the historical trajectory of the green robot, which exhibits relatively lower uncertainty. The construction of a inter-robot virtual map for map uncertainty estimation could potentially lead to conflicts in decision-making for a local robot. As outlined in Alg. 2, we generate a local virtual map denoted as $\mathcal{V}_\alpha^{\text{new}*}$ for robot α. This map is constructed using $\mathcal{X}_\alpha^{\text{new}*}$ and corresponding observations $\mathcal{Z}_\alpha^{\text{new}}$. Then, we compute the uncertainty utility of the map by considering the overlapping observed segment between the current inter-robot virtual map $\mathcal{V}^*$ and the predicted local virtual map $\mathcal{V}_\alpha^{\text{new}*}$ with Eq. (9).

D. Experiments and Results

We perform experiments within a 200m × 200m environment, employing randomly generated landmarks. The errors described below define 95% confidence intervals. We assume each robot is furnished with a sonar with range error 0.1m, bearing error 0.1°, and max. sensing range 10m. Each robot also performs inertial dead reckoning; its gyro and accelerometer yield errors of 5° and 0.05m. To navigate towards uncharted territories, we utilize the Artificial Potential Field (APF) method, detailed in [14], to avoid collision. The result of an environment exploration involving five robots is depicted in Fig. 3.

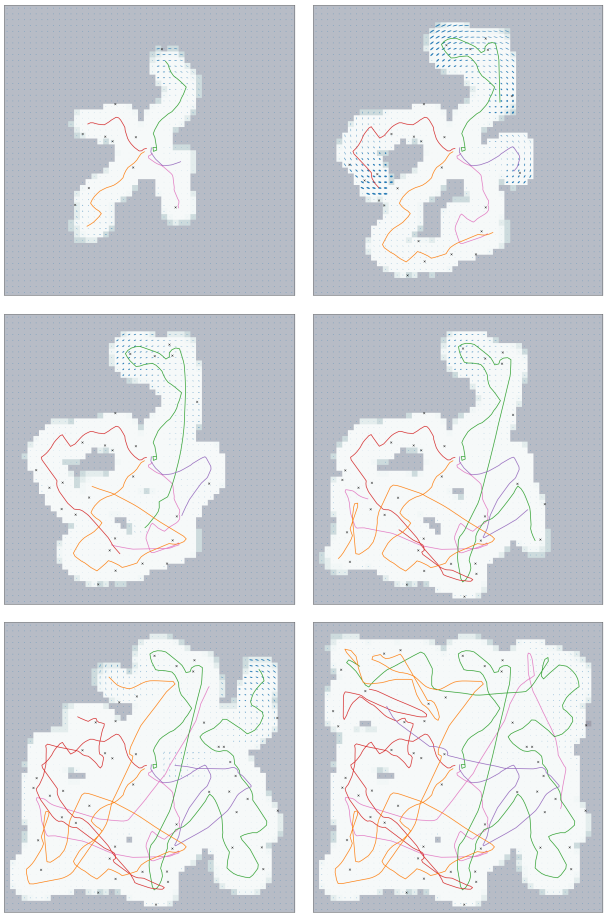

Fig. 3: An example of five robots exploring a $200m \times 200m$ environment with the proposed algorithm.

References

[1] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023.

[2] 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.

[3] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on robotics*, vol. 21, no. 3, pp. 376–386, 2005.

[4] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart, "Distributed multirobot exploration and mapping," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1325–1339, 2006.

[5] D. Jang, J. Yoo, C. Y. Son, D. Kim, and H. J. Kim, "Multi-robot active sensing and environmental model learning with distributed gaussian process," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5905–5912, 2020.

[6] J. Yu, J. Tong, Y. Xu, Z. Xu, H. Dong, T. Yang, and Y. Wang, "Smmr-explore: Submap-based multi-robot exploration system with multi-robot multi-target potential field exploration method," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8779–8785.

[7] N. Atanasov, J. Le Ny, K. Daniilidis, and G. J. Pappas, "Decentralized active information acquisition: Theory and application to multi-robot slam," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 4775–4782.

[8] M. Kontitsis, E. A. Theodorou, and E. Todorov, "Multi-robot active slam with relative entropy optimization," in *2013 American Control Conference*. IEEE, 2013, pp. 2757–2764.

[9] M. Corah, C. O'Meadhra, K. Goel, and N. Michael, "Communication-efficient planning and mapping for multi-robot exploration in large environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1715–1721, 2019.

[10] J. Wang and B. Englot, "Autonomous exploration with expectation-maximization," in *Robotics Research: The 18th International Symposium ISRR*. Springer, 2020, pp. 759–774.

[11] 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.

[12] J. Wang, F. Chen, Y. Huang, J. McConnell, T. Shan, and B. Englot, "Virtual maps for autonomous exploration of cluttered underwater environments," *IEEE Journal of Oceanic Engineering*, vol. 47, no. 4, pp. 916–935, 2022.

[13] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.

[14] X. Fan, Y. Guo, H. Liu, B. Wei, and W. Lyu, "Improved artificial potential field method applied for AUV path planning," *Mathematical Problems in Engineering*, vol. 2020, pp. 1–21, 2020.