

Real-Time Planning Under Uncertainty for AUVs Using Virtual Maps

Ivana Collado-Gonzalez, John McConnell, Jinkun Wang, Paul Szenher, and Brendan Englot

Abstract—Reliable localization is an essential capability for marine robots navigating in GPS-denied environments. SLAM, commonly used to mitigate dead reckoning errors, still fails in feature-sparse environments or with limited-range sensors. Pose estimation can be improved by incorporating the uncertainty prediction of future poses into the planning process and choosing actions that reduce uncertainty. However, performing belief propagation is computationally costly, especially when operating in large-scale environments. This work proposes a computationally efficient planning under uncertainty framework suitable for large-scale, feature-sparse environments. Our strategy leverages SLAM graph and occupancy map data obtained from a prior exploration phase to create a *virtual map*, describing the uncertainty of each map cell using a multivariate Gaussian. The virtual map is then used as a cost map in the planning phase, and performing belief propagation at each step is avoided. A receding horizon planning strategy is implemented, managing a goal-reaching and uncertainty-reduction tradeoff. Simulation experiments in a realistic underwater environment validate this approach. Experimental comparisons against a full belief propagation approach and a standard shortest-distance approach are conducted.

I. INTRODUCTION

Autonomous Underwater Vehicles (AUVs) are being increasingly used to perform tasks such as inspection, mapping, and exploration. Since GPS is unavailable underwater, AUVs must rely on onboard sensors to estimate their position, making them susceptible to drift. Simultaneous Localization and Mapping (SLAM) has been successfully applied to correct dead reckoning drift. However, SLAM requires environmental features to lie within the AUV's field of view. Therefore, robots operating in feature-sparse environments or with short-range sensors are especially prone to accumulating pose estimation errors. Thus, executing trajectories that do not account for robot sensing capabilities leads to high localization uncertainty. Planning under uncertainty, however, commonly relies on belief propagation, which is computationally expensive and can become infeasible to compute when operating in large-scale environments.

Recent work has focused on reducing map and robot pose uncertainty during the exploration of unknown environments, without considering navigation tasks that follow the mapping mission [1]. Furthermore, strategies that do consider goal-reaching tasks commonly assume no prior knowledge [2] or perfect prior knowledge of the environment. However, exploration is an imperfect process, and the resulting maps are corrupted by noise, uncertainty, and incomplete information, and should be treated accordingly when used

I. Collado-Gonzalez, P. Szenher, and B. Englot are with Stevens Institute of Technology, Hoboken, NJ, USA, {icollado, pszenher, benglot}@stevens.edu. J. McConnell was previously with Stevens and is now with the United States Naval Academy, Annapolis, MD, USA, jmcconne@usna.edu. J. Wang was previously with Stevens and is now with Anyware Robotics, Fremont, CA USA. **This research was supported by NSF Grant IIS-1652064 and USDA-NIFA Grant 2021-67022-35977.**




Fig. 1: **Planning under uncertainty example.** **Top:** Gazebo environment and planning example. Blue star is start location, red cross shows goal location, pink arrows show candidate place-revisiting actions, and green diamond represents goal-reaching candidate action. **Bottom:** Virtual map and representative resulting paths using different uncertainty handling strategies. *Shortest Distance* is shown in red, *Expectation Maximization* in blue, and our proposed strategy in green. The uncertainty growth from executing each path is illustrated along the path using 95% confidence ellipses.

for subsequent planning queries. Planning to a specified goal can similarly be corrupted by noise and uncertainty. Choosing an uninformative trajectory to reach a goal can lead to substantial uncertainty growth, causing robots to crash or become "lost". Instead of discarding SLAM graph information after exploration, this information can be used to assess the uncertainty of the resulting map, informing any subsequent path planning queries. Moreover, by performing SLAM during navigation, we can adapt to small changes in the environment, manage errors from noise, and deal with incomplete information.

In this paper, we leverage mapping and SLAM graph information from a prior exploration run and use it to create a *virtual map* that will in turn be used to perform efficient planning under uncertainty, without the need for computationally demanding online belief propagation. Specifically, we present a planning under uncertainty scheme for underwater robots operating at a fixed depth in large-scale feature-sparse environments. Our main contributions are as follows:

- To our knowledge, the first algorithmic pipeline for underwater robots that uses exploration as a pre-processing step for planning under uncertainty to specified goals.
- We leverage the resulting SLAM graph and occupancy map information to produce a useful costmap for computationally efficient planning under uncertainty.
- Realistic simulation experiments, in which receding horizon planning is used to compare our proposed

method with shortest-distance and full belief propagation planning strategies.

The rest of this work is structured as follows. First, Sec. II describes related works. Next, the problem formulation is explained in Sec. III, followed by the proposed method in Sec. IV. Subsequently, results and experiments are presented in Sec. V. Finally, in Sec. VI conclusions are drawn.

II. RELATED WORKS

Incorporating localization uncertainty into the planning process is important for navigation tasks including exploration and goal-reaching missions. Actively maintaining high localization accuracy while navigating was first addressed in [3], by controlling robot motion and sensing actions. Many mobile robots, including AUVs, rely heavily on perceptual features to localize in their environment, thus perception has a substantial impact on pose uncertainty and has been studied in several works [4], [5], [6]. Moreover, in the context of SLAM, revisiting previously seen features is necessary to trigger loop closures and lower uncertainty. Image saliency was used in underwater active SLAM [7] to identify uncertainty reducing actions. Some methods [8], [9] attempt to reduce uncertainty by performing rotational actions to observe visual features while following a path. However, only rotational movements are not sufficient to localize when features are out of sensing range.

Therefore, balancing the trade-off between reducing uncertainty and completing an exploration or goal-reaching task is an important part of planning under uncertainty strategies. To limit deviations from the main path, [7] and [9] only compute and perform uncertainty reducing actions when the uncertainty of the robot pose is above a certain threshold. Most active SLAM motion planning approaches employ three stages: goal identification, utility computation, and action selection [10]. Furthermore, the receding horizon path planning strategy, in which robots execute only a small portion of each plan prior to replanning, has been used to solve the active SLAM problem in numerous works [6], [7], where the continuous feedback of mapping the environment helps attenuate errors in tracking and perception.

Furthermore, planning under uncertainty has been tackled using Partially Observable Markov Decision Process (POMDP)s [11]. However, POMDP approaches grow exponentially in complexity as the number of actions and observations increases. To reduce computational burden various strategies have been proposed. The Belief Roadmap (BRM) [12] made belief space planning more efficient by factoring the covariance matrix and combining multiple EKF update steps. Still, it results in significant computation time for running on a complex graph in practical problems. Rapidly-exploring Random Belief Trees (RRBTs) [13] extend the fast computing Rapidly-exploring Random Tree (RRT)* [14] framework to take pose uncertainty into account.

In the above methods, performing perception-related belief propagation is crucial, but often the computational bottleneck. Hence, another body of work has focused on developing dedicated representations of the environment and

the uncertainty for efficient computation. A coarse-resolution map for rapidly evaluating the information gain of planned actions was used in [15]. Similarly, working with a reduced local map instead of a full map was done in [16]. To avoid online belief propagation, [17] pre-computes and saves the information content of the environment in a 2D grid which can be used during the robot's mission. [18] proposes a localizability matrix over a probability grid map which includes localizability index and direction. Yet, this approach only considers active localization and does not perform path planning. In [19], the most feature-rich path for a rover is calculated by analyzing the soil map of Mars, assuming a constant covariance accumulation rate in feature-rich and feature-sparse areas. A technique that turns the information from publicly available maps into a localizability map is presented in [20]. This approach uses graph SLAM, but does not incorporate path planning.

Expectation-Maximization (EM) exploration [21], [22], [23] presents the concept of a *virtual map*, in which every map cell's uncertainty is represented by a multivariate Gaussian, and the map uncertainty resulting from future actions can be predicted. In our present work, we do not focus on exploration, but on the subsequent planning to specified goal states after exploration has been performed. We adapt the virtual map presented in [21], and use it as an uncertainty cost map for computationally efficient planning, avoiding the need for online belief propagation. Additionally, we use a receding horizon path planning strategy that manages a goal-reaching and place-revisiting trade-off. In the sections below, we will describe our proposed planning under uncertainty framework and its application to underwater navigation with sonar, demonstrating that it achieves far lower localization uncertainty and error than a standard shortest-distance approach, and that the errors obtained are comparable to that of full-fledged EM belief propagation.

III. BACKGROUND AND PROBLEM FORMULATION

A. Simultaneous Localization And Mapping

We use an incremental graph-based pose SLAM approach. Thus, the entire robot trajectory is repeatedly estimated, allowing the robot pose history and map to be corrected throughout the mission. The motion model is defined by

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

where $\mathbf{x}_i = [x_i, y_i, \theta_i]$ is the state vector describing position and orientation in $\text{SE}(2)$, $\mathbf{u}_i$ is the control input vector, and $\mathbf{w}_i$ is additive process noise with covariance matrix $\mathbf{Q}_i$. The sensor model is given by

$$\mathbf{z}_k = h_k(\mathbf{x}_{i_k}) + \mathbf{q}_k, \quad \mathbf{q}_k \sim \mathcal{N}(0, \mathbf{R}_k), \quad (2)$$

where we obtain measurements $\mathbf{z}_k$ at pose $\mathbf{x}_{i_k}$, corrupted by additive sensor noise $\mathbf{q}_k$ with covariance matrix $\mathbf{R}_k$.

Let $\mathcal{X} = \{\mathbf{x}_{i=0}^T\}$ be the set of poses from time 0 to time T . $\mathcal{C}$ contains all constraints between robot poses. For each pair of poses defining a constraint $\langle i, j \rangle \in \mathcal{C}$, the

error $\mathbf{e}_{ij}$ between observed transformation $\mathbf{z}_{ij}$ and predicted transformation $\hat{\mathbf{z}}_{ij}$ is defined as:

$$\mathbf{e}_{ij}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{z}_{ij} - \hat{\mathbf{z}}_{ij}(\mathbf{x}_i, \mathbf{x}_j), \quad (3)$$

$$\hat{\mathbf{z}}_{ij}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^\top \mathbf{x}_j. \quad (4)$$

The SLAM problem can be represented as a nonlinear least squares problem [24]:

$$\mathcal{X}^* = \arg \min_{\mathbf{x}} \sum_{(i,j) \in \mathcal{C}} \mathbf{F}_{ij}, \quad \mathbf{F}_{ij} = \mathbf{e}_{ij}^\top \boldsymbol{\Omega}_{ij} \mathbf{e}_{ij}, \quad (5)$$

where $\mathbf{F}_{ij}$ denotes the negative log-likelihood function of one constraint between $\mathbf{x}_i$ and $\mathbf{x}_j$, and $\boldsymbol{\Omega}_{ij}$ is the information matrix. The SLAM system aims to find a set of robot poses which minimize the total observation error.

B. Virtual Map

Unlike occupancy grid maps, the virtual map [23] encodes map cell uncertainty (represented by a multivariate Gaussian) instead of occupancy probability. Let the virtual map $\mathcal{V} = \{\mathbf{v}_i\}$ be the set of possible observations located at each map cell $\mathbf{v}_i \in \mathbb{R}^2$, distributed as $\mathcal{N}(\mathbf{v}_i, \Sigma_{\mathbf{v}_i})$, where $\mathbf{v}_i = [x_i, y_i]$ is the cell location and $\Sigma_{\mathbf{v}_i}$ is the corresponding covariance matrix. When initializing a virtual map, a prior with conservatively high initial uncertainty is imposed on all map cells.

The map cell covariances are updated based on the current estimate of the trajectory $\mathcal{X}_{\text{old}}$ and the history of measurements $\mathcal{Z}$.

$$q(\mathcal{V}) = p(\mathcal{V} | \mathcal{X}_{\text{old}}, \mathcal{Z}) \quad (6)$$

The maximum posterior probability estimate for the virtual map can thus be computed by

$$\mathcal{V}^* = \arg \max_{\mathcal{V}} p(\mathcal{V} | \mathcal{X}_{\text{old}}, \mathcal{Z}). \quad (7)$$

It is assumed that measurements are assigned to maximize the likelihood

$$\mathcal{Z} = \arg \max_{\mathcal{Z}} h(\mathcal{X}, \mathcal{V}). \quad (8)$$

By measuring environmental features, we can derive the transformations between poses with overlapping observations using Iterative Closest Point (ICP). We essentially aim to minimize the closeness of measurements using ICP-based pose SLAM. For more details regarding virtual map pose covariance computation, the reader should consult [23].

C. Problem Statement

In this paper, the problem of autonomously navigating to a goal location under pose uncertainty is addressed. We adapt the virtual map, previously used for exploration, to accommodate planning under uncertainty to specified goals. We leverage SLAM information from a prior exploration run to compute uncertainty values offline and store them in the form of a virtual map $\mathcal{V}$. We then focus on finding a path to a goal that curbs the growth of uncertainty using pose SLAM and the covariance information from our virtual map. Specifically, the virtual map is used as a cost map during the planning stage to estimate future uncertainty, and the SLAM graph is maintained and updated to help mitigate dead reckoning errors and to update our occupancy map used for planning, as the prior map is imperfect.




Fig. 2: **System overview:** SLAM information from a prior exploration run is used to initialize the current SLAM instance and build a virtual map offline. The motion planner takes the accompanying occupancy map and searches for candidate paths. The utility of each candidate path is assessed using the virtual map. The chosen path and the current robot pose are then passed to the guidance module, which provides the desired velocity vector to follow the desired path. The low-level PID controller takes the desired velocity values and defines the required thruster forces.

IV. METHODOLOGY

In this section, we explain our proposed framework, the system overview can be observed in Fig. 2. First, SLAM information from a prior exploration run is taken and used to initialize our current SLAM instance and build our virtual map $\mathcal{V}$. The motion planner takes the occupancy map $\mathcal{M}$ from SLAM and searches for candidate paths $\mathcal{P} = \{\mathcal{X}_c\}$. Then, the utility of each candidate path is assessed using the virtual map. The chosen path $\mathcal{X}_d$ and the current robot pose $\mathbf{x}$ are then passed to the pure pursuit guidance module, which in turn provides the appropriate desired velocity vector $\mathbf{u}$ to follow the desired path. The low-level PID controller takes the desired velocity values and defines thruster forces τ .

A. Sonar SLAM

To estimate the location of the robot, we use the same graph-based pose SLAM framework as in [23]. The sensory inputs include an Inertial Measurement Unit (IMU), a Doppler Velocity Log (DVL), and a sonar image. The sonar image is first processed using Constant False Alarm Rate (CFAR) detection [25] to identify returns from the surrounding environment. All CFAR-extracted features in the image are converted to Cartesian coordinates. The SLAM factor graph is denoted as

$$\mathbf{f}(\boldsymbol{\Theta}) = \mathbf{f}^0(\boldsymbol{\Theta}_0) \prod_i \mathbf{f}_i^0(\boldsymbol{\Theta}_i) \prod_j \mathbf{f}_j^{\text{LC}}(\boldsymbol{\Theta}_j),$$

where variables $\boldsymbol{\Theta}$ contain 3-DOF robot poses and $\mathbf{f}^0$ denotes the initial factor graph. The odometry factor $\mathbf{f}^0$ defines the relative motion between two consecutive poses from persistent odometry measurements obtained by DVL/IMU dead reckoning. Loop closure factors $\mathbf{f}^{\text{LC}}$ are obtained by identifying non-sequential scan matches, applying Iterative Closest Point (ICP) [26] between the current sonar keyframe and frames inside a given search radius. ICP is initialized using a DVL/IMU dead reckoning pose estimate and further optimized using consensus set maximization [27], helping ICP to avoid local minima. Loop closure outliers are rejected by evaluating point cloud overlap and applying Pairwise Consistent Measurement Set Maximization (PCM) [28]. At the back end, we employ the GTSAM [29] implementation of iSAM2 [30], which is merged with dead reckoning to provide a high-frequency state estimate. Occupancy map $\mathcal{M}$ is updated and used for path planning.

Input: Occupied cells $\mathcal{O}$, number of place-revisiting actions N_{pr} , revisiting radius r_{pr} , minimum separation distance d_{pr} , shortest path to goal $\mathcal{X}_{\text{sd}}$, number of clusters k

Output: $\mathcal{A}_{\text{pr}}, \mathcal{R} \leftarrow \emptyset$

$\mathcal{M} = \{\mathbf{m}_1, \dots, \mathbf{m}_k\} \leftarrow \text{ComputeKMeans}(\mathcal{O}, k)$

for $i \leftarrow 1$ **to** k **do**

$\mathcal{C} \leftarrow \{(x_i + r_{\text{pr}} \cos(\theta), y_i + r_{\text{pr}} \sin(\theta), \theta) \mid \theta \in [0, 2\pi)\}$

$\mathcal{R} \leftarrow \mathcal{R} \cup \{\mathcal{C}\}$

while $|\mathcal{A}_{\text{pr}}| < N_{\text{pr}}$ **do**

$\mathbf{a}^* \leftarrow \arg \min_{\mathbf{r}_i \in \mathcal{R}} \|\mathbf{r}_i - \mathcal{X}_{\text{sd}}\|^2$

$d^* \leftarrow \min_{\mathbf{g}_i \in \mathcal{A}_{\text{pr}}} \|\mathbf{a}^* - \mathbf{g}_i\|$

if $d^* \leq d_{\text{pr}}$ **then**

$\mathcal{A}_{\text{pr}} \leftarrow \mathcal{A}_{\text{pr}} \cup \{\mathbf{a}^*\}$

$\mathcal{R} \leftarrow \mathcal{R} \setminus \{\mathbf{a}^*\}$

return Candidate place-revisiting actions $\mathcal{A}_{\text{pr}}$

B. Motion Planning

Global trajectory optimization in a large-scale planning problem is computationally expensive, and may require replanning as the trajectory is executed. Therefore, we apply a receding horizon path planner, which also performs trajectory tracking and incorporates perceptual feedback. Inspired by many active SLAM motion planning approaches [10], we employ three stages: goal identification, utility computation, and action selection. Action selection and path planning are treated as tightly coupled problems. We consider two different action types, one type for lowering uncertainty (place-revisiting), and another for reaching the mission goal (shortest-distance). *Place-revisiting* actions look for poses that can lower uncertainty by revisiting locations the robot has previously observed, triggering SLAM loop closures. *Shortest-distance* actions represent locations that lay on the shortest feasible path to the goal. Fig. 1a depicts place-revisiting actions as pink arrows and the shortest-distance action as a green diamond.

First, a finely discretized roadmap is generated at the beginning of the mission. The boundaries of the roadmap are defined, and edges are pruned out if obstacles are discovered. The A* algorithm [31] is used to find the shortest path from the start position to the final mission goal without colliding with known obstacles. Afterward, the shortest-distance action set $\mathcal{A}_{\text{sd}}$ is defined by choosing targets at user-defined distances d_{sd} from the current position along the identified shortest distance path. For safety, the distance to the waypoints should be less or equal to the robot sensing range to avoid collision with possible new obstacles and refrain from unnecessary replanning. Likewise, traversing large featureless distances should be prevented to avoid large uncertainty growth.

The place-revisiting action set $\mathcal{A}_{\text{pr}}$ is defined following Algorithm 1. The target locations are chosen by identifying action candidates that are close to the shortest path while observing occupied cells. These poses have the potential of triggering SLAM loop closures while avoiding large deviations from the desired path. First, occupied cells in

the occupancy map are identified and clustered using the k-means algorithm. Afterward, candidate goals are sampled along the boundary of a circle centered at each cluster origin with a radius r_{pr} that is user-specified. We then choose the candidates that are closest to the shortest path. Finally, the A* search algorithm is used to find paths from the current position to the final goal, each path passing through one place-revisiting waypoint.

From the candidate paths $\mathcal{P}$ identified, one is chosen using the utility function defined in Sec. IV-C. Once the desired path $\mathcal{X}_{\text{d}}$ is identified, and the vehicle arrives at the place-revisiting or shortest-distance target location, the planning strategy is repeated until the robot arrives at its final goal.

C. Utility evaluation

We define the utility function used to choose the desired path $\mathcal{X}_{\text{d}} = \arg \max_{X \in \mathcal{P}} U(X_{T:T+N})$ as follows:

$$U(X_{T:T+N}) = - \sum_{i=1}^n \log \det(\Sigma_{\mathbf{v}_i}) - \alpha d(X_{T:T+N}). \quad (9)$$

The uncertainty metric used here is the sum of the log-determinant of the covariance matrices of the virtual map cells $\Sigma_{\mathbf{v}_i}$ that will be encountered by the chosen path. The entire path from the current position to the target position is considered. In addition to the uncertainty term, a distance cost weighted by α is added to encourage a trade-off between traveling and uncertainty reduction. A path will be discarded and recomputed if it is found to be blocked by an obstacle.

D. Virtual Map Adaptation

In contrast to the original implementation of virtual maps, we compute virtual map values only once at the start of a planning run. First, the SLAM graph and occupancy map information from a prior SLAM exploration run are loaded. Then, the past trajectory information $\mathcal{X}_{\text{old}}$ is used to estimate virtual map covariances $\Sigma_{\mathbf{v}_i}$. Even though we can use the same map resolution for path planning and uncertainty estimation, it was shown in [21] that a lower resolution virtual map provides similar performance but permits faster belief propagation. In the interest of computational efficiency and fair experimental comparisons, we discretize the virtual map to be as coarse as possible, while still capturing the essential geometric features of the robot's workspace.

V. EXPERIMENTS AND RESULTS

A. Experiment Setup

To validate the proposed planning under uncertainty framework, we present experimental results from a high-fidelity AUV simulation operating at a fixed depth in different environments. The simulation setup ran on a Lambda Workstation with an AMD 3970X 32-core processing unit and 256GB RAM, running Ubuntu 18.04, ROS Melodic, and Gazebo (our planning implementation is serialized and every trial used only a single core). The vehicle used in the simulations is the RexROV2 from UUV Simulator [32]. The simulated sonar operates at 5 Hz, has 30 meter range,




Fig. 3: **Virtual maps for all three simulation environments:** including planning start (blue star) and all goal positions (numbered red x's).

and has a horizontal aperture of $\theta = [-65^\circ, 65^\circ]$. Zero-mean Gaussian noise is added to range and bearing sonar measurements: $\sigma_r = 0.2$ m, $\sigma_\theta = 0.02$ rad. Linear velocity values are obtained from a simulated Doppler Velocity Log (DVL) and angular measurements are obtained from a simulated Inertial Measurement Unit (IMU) at 50 Hz and 7 Hz respectively. Likewise, zero-mean Gaussian noise is added with $\sigma_x = \sigma_y = 0.1$ m/s and $\sigma_\theta = 0.015$ rad.

Three environments were designed to simulate realistic large-scale scenarios where an AUV might operate. The first environment, a marina, was inspired by Penn's Landing in Philadelphia, the second scenario is an offshore fish farm containing six cylindrical fish pens, and the third scenario is a bridge-tunnel system, inspired by the Chesapeake Bay bridge-tunnel. As a preprocessing step, each environment was fully explored while performing sonar SLAM as presented in Sec. IV-A [23]. Our AUV was teleoperated inside its simulated environment to efficiently complete the exploration task, during which a globally deformable occupancy grid map was constructed based on local submaps [33]. This representation anchors a local submap to each SLAM pose keyframe, allowing for efficient map recomputation when a trajectory segment is changed. The trajectories used for each exploration can be observed in the video attachment. The SLAM graph and occupancy map information were saved and later reloaded to build the virtual map and use our planner. The occupancy grid map was downsampled from 0.2m to 2m in cell size to produce the virtual map.

Each scenario has one start location and three different goal locations. Each run consisted of the robot navigating from a start to a goal location, and was repeated for 100 trials. The results shown are the average values of all the trials executed. Additionally, we compare our planner




Fig. 4: **Marina results:** goals 1, 2, 3 shown from top to bottom.

framework with two other strategies, Shortest Distance (SD) and Expectation-Maximization (EM). The *SD planner* just follows the shortest viable A* path to the goal. The *EM planner* uses the same receding horizon style from Sec. IV-B, however, it employs full Expectation-Maximization belief propagation as in [23] to choose each intermediate goal.

For our experiments, we chose to have one shortest-distance action $N_{sd} = 1$ at $d_{sd} = 25$ m distance, and five place revisiting actions $N_{pr} = 5$ each iteration. The place revisiting radius was defined at $r_{pr} = 15$ m. The distance weight was experimentally chosen and set to $\alpha = 3$. To evaluate our planning framework, we recorded the pose uncertainty of the current robot pose, computed as $\det(\Sigma_{x_i})^{\frac{1}{3}}$, and the pose error, computed via RMSE, across all trials as functions of distance traveled.

B. Results Comparison

In our results, SD refers to the *Shortest Distance* planner, EM refers to the *Expectation-Maximization* planner, and VM refers to our proposed planner, which uses a virtual map as a costmap to approximate belief propagation.

1) *Marina Environment:* The virtual map for the marina environment can be observed in Fig. 3a, where the blue star represents the start location, and the red crosses depict goal locations. Pose error and uncertainty results for each goal location can be observed in Fig. 4. VM and EM planners achieve a similarly low uncertainty performance, while SD uncertainty grows unbounded. Goal 3 results in Fig. 4 (bottom) show VM and EM manage to reduce the initial robot pose error, while SD does not. Surprisingly, goal 1 results in Fig. 4 (top) reveal VM achieves a more stable low pose error performance than EM, likely due to EM having an erroneous loop closure. Goal 2 results in Fig. 4 (middle) convey that EM has the best pose error result,




Fig. 5: **Fish farm results:** goals 1, 2, 3 shown from top to bottom.

while VM has intermediate performance, still better than SD. Here, VM likely missed an opportunity to lower pose error by not predicting and achieving a SLAM loop closure.

2) *Fish Farm Environment:* The virtual map for the fish farm environment with start (blue star) and goal locations (red cross) can be seen in Fig. 3b, and corresponding results can be observed in Fig. 5. Resembling the previous experiment, VM and EM achieve a similarly low uncertainty performance while SD uncertainty has a higher and less stable uncertainty result. Fig. 5 (top) shows that for goal 1 the SD planner manages to reduce its pose error, before pose error spikes up again. Likewise, Fig. 5 (bottom) shows that for goal 3 the SD planner manages to lower the pose error after a spike, potentially due to a coincidental loop closure. However, EM and VM still achieve lower pose error than SD in all three scenarios, with EM being slightly superior.

3) *Bridge-tunnel Environment:* The virtual map and relevant locations for the bridge-tunnel environment are shown in Fig. 3c, and results are shown in Fig. 6. For all three targets, VM performance is similar to EM, while SD uncertainty grows unbounded and exhibits higher pose estimation errors.

4) *Result Summary:* Fig. 1 shows example paths achieved by each planner, where the ellipses represent the evolution of uncertainty. It is clear that the VM and EM paths are similar in error and uncertainty; they both keep close to features in the environment and achieve low uncertainty by triggering SLAM loop closures. VM and EM perform drastically better lowering uncertainty than the SD planner. Lastly, Table I shows the amount of time it takes to evaluate each candidate action and choose the appropriate desired path. In the SD planner, only one path is assessed and chosen. In our proposed VM planner and in the EM planner, 6 candidate actions are assessed each iteration. The results show the average time of every utility evaluation performed




Fig. 6: **Bridge-tunnel results:** goals 1, 2, 3 shown from top to bottom.

	Computation Time (sec)					
	EM		VM		SD	
	(μ)	(σ)	(μ)	(σ)	(μ)	(σ)
Marina 1	3.3553	0.1446	0.0041	0.0001	0.0015	0.0002
Marina 2	2.9324	0.0598	0.0039	0.0002	0.0014	0.0002
Marina 3	3.1507	0.1396	0.0040	0.0001	0.0015	0.0004
Fish Farm 1	5.4149	0.2094	0.0057	0.0001	0.0023	0.0001
Fish Farm 2	4.9547	0.1768	0.0054	0.0001	0.0021	0.0001
Fish Farm 3	5.4370	0.2138	0.0057	0.0001	0.0023	0.0001
Bridge-Tunnel 1	6.2699	0.0744	0.0069	0.0001	0.0027	0.0001
Bridge-Tunnel 2	6.4998	0.1647	0.0070	0.0001	0.0027	0.0001
Bridge-Tunnel 3	6.1624	0.1358	0.0067	0.0001	0.0026	0.0002

TABLE I: **Computation time comparison:** we give the mean and standard deviation of the utility evaluation time required in each experiment.

in all 100 trials. As expected, SD planner has the fastest computation time while EM has the slowest. It can be seen that EM takes 3 orders of magnitude more time than VM, while VM and SD maintain the same order of magnitude.

VI. CONCLUSION

This paper presents a computationally efficient planning under uncertainty framework for underwater robots operating in large-scale feature-sparse environments. We adapt the concept of virtual maps, a product of a robot's initial exploration of a new environment, using them as costmaps to avoid expensive online belief propagation while planning. A receding horizon motion planning strategy is implemented for planning to specified goals in a manner that is compatible with real AUV operations. The advantages of the proposed framework are demonstrated in a realistic underwater simulation. Results show a decrease in uncertainty and pose error compared to a standard shortest-distance approach. Furthermore, our approach is much faster than full belief propagation, while still maintaining low uncertainty and pose error. We aim for future work to extend this strategy to 3D planning scenarios, and to perform virtual map updates to accommodate important changes in the environment.

REFERENCES

- [1] 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.
- [2] E. Pairet, J. D. Hernández, M. Carreras, Y. Petillot, and M. Lahijanian, "Online mapping and motion planning under uncertainty for safe navigation in unknown environments," *IEEE Transactions on Automation Science and Engineering*, vol. 19, no. 4, pp. 3356–3378, 2022.
- [3] W. Burgard, D. Fox, and S. Thrun, "Active mobile robot localization," *International Joint Conference on Artificial Intelligence*, 1997.
- [4] C. Papachristos, S. Khattak, and K. Alexis, "Uncertainty-aware receding horizon exploration and mapping using aerial robots," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4568–4575, 2017.
- [5] G. Costante, C. Forster, J. Delmerico, P. Valigi, and D. Scaramuzza, "Perception-aware path planning," *arXiv preprint arXiv:1605.04151*, 2016.
- [6] Z. Zhang and D. Scaramuzza, "Perception-aware receding horizon navigation for MAVs," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2534–2541, 2018.
- [7] S. Suresh, P. Sodhi, J. G. Mangelson, D. Wettergreen, and M. Kaess, "Active SLAM using 3D submap saliency for underwater volumetric exploration," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3132–3138, 2020.
- [8] I. Maurović, M. Seder, K. Lenac, and I. Petrović, "Path planning for active SLAM based on the D* algorithm with negative edge weights," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 48, no. 8, pp. 1321–1331, 2018.
- [9] F. Andrade, M. Llofriu, M. Marzoa Tanco, G. Trinidad Barnech, and G. Tejera, "Active localization strategy for hypotheses pruning in challenging environments," *Journal of Intelligent & Robotic Systems*, 2022.
- [10] 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.
- [11] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, no. 1, pp. 99–134, 1998.
- [12] S. Prentice and N. Roy, "The belief roadmap: Efficient planning in belief space by factoring the covariance," *The International Journal of Robotics Research*, vol. 28, pp. 1448 – 1465, 2009.
- [13] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 723–730, 2011.
- [14] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [15] M. Nealson, Erik Corah and N. Michael, "Environment model adaptation for mobile robot exploration," *Autonomous Robots*, vol. 42, 2018.
- [16] F. Andrade, M. Llofriu, M. M. Tanco, G. T. Barnech, and G. Tejera, "Active localization for mobile service robots in symmetrical and open environments," *Latin American Robotics Symposium (LARS)*, pp. 270–275, 2021.
- [17] N. Roy, W. Burgard, D. Fox, and S. Thrun, "Coastal navigation-mobile robot navigation with uncertainty in dynamic environments," *IEEE International Conference on Robotics and Automation (ICRA)*, vol. 1, pp. 35–40 vol.1, 1999.
- [18] Z. Liu, W. Chen, Y. Wang, and J. Wang, "Localizability estimation for mobile robots based on probabilistic grid map and its applications to localization," *IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pp. 46–51, 2012.
- [19] H. Inoue, M. Ono, S. Tamaki, and S. Adachi, "Active localization for planetary rovers," *IEEE Aerospace Conference*, pp. 1–7, 2016.
- [20] V. Olga and S. Cyrill, "Improving SLAM by exploiting building information from publicly available maps and localization priors," *Journal of Photogrammetry, Remote Sensing and Geoinformation Science*, vol. 85, 2017.
- [21] J. Wang and B. Englot, "Autonomous exploration with expectation-maximization," *International Symposium on Robotics Research (ISRR)*, pp. 759–774, 2017.
- [22] J. Wang, T. Shan, and B. Englot, "Virtual maps for autonomous exploration with pose SLAM," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4899–4906, 2019.
- [23] 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.
- [24] G. Grisetti, R. Kümmerle, C. Stachniss, and W. Burgard, "A tutorial on graph-based SLAM," *IEEE Intelligent Transportation Systems Magazine*, vol. 2, no. 4, pp. 31–43, 2010.
- [25] M. Richards, "Fundamentals of radar signal processing," *McGraw-Hill Education*, 2005.
- [26] P. 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.
- [27] M. A. Fischler and R. C. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Readings in Computer Vision*, pp. 726–740, 1987.
- [28] J. G. Mangelson, D. Dominic, R. M. Eustice, and R. Vasudevan, "Pairwise consistent measurement set maximization for robust multi-robot map merging," *IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2916–2923, 2018.
- [29] F. Dellaert, "Factor graphs and GTSAM: A hands-on introduction," *Georgia Institute of Technology*, Tech. Rep. GT-RIM-CP&R-2012-002, 2012.
- [30] 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.
- [31] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [32] M. M. M. Manhães, S. A. Scherer, M. Voss, L. R. Douat, and T. Rauschenbach, "UUV simulator: A gazebo-based package for underwater intervention and multi-robot simulation," *OCEANS MTS/IEEE Monterey*, pp. 1–8, 2016.
- [33] B.-J. Ho, P. Sodhi, P. Teixeira, M. Hsiao, T. Kusnur, and M. Kaess, "Virtual occupancy grid map for submap-based pose graph SLAM and planning in 3D environments," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2175–2182, 2018.