Informative Path Planning for Active Field Mapping under Localization Uncertainty

Marija Popović[1], Teresa Vidal-Calleja[2], Jen Jen Chung[3], Juan Nieto[3], Roland Siegwart[3]

***Abstract*— Information gathering algorithms play a key role in unlocking the potential of robots for efficient data collection in a wide range of applications. However, most existing strategies neglect the fundamental problem of the robot pose uncertainty, which is an implicit requirement for creating robust, high-quality maps. To address this issue, we introduce an informative planning framework for active mapping that explicitly accounts for the pose uncertainty in *both* the mapping and planning tasks. Our strategy exploits a Gaussian Process (GP) model to capture a target environmental field given the uncertainty on its inputs. For planning, we formulate a new utility function that couples the localization and field mapping objectives in GP-based mapping scenarios in a principled way, without relying on manually-tuned parameters. Extensive simulations show that our approach outperforms existing strategies, reducing mean pose uncertainty and map error. We present a proof of concept in an indoor temperature mapping scenario.**

I. INTRODUCTION

Rapid technological advancements are inciting the use of autonomous mobile robots for exploration and data acquisition. In many marine [1, 2], terrestrial [3, 4], and airborne [5, 6] applications, these systems have the ability to bridge the spatiotemporal divides limiting traditional measurement methods in a safer and more cost-effective manner [7]. However, to fully exploit their potential, algorithms are required for planning efficient informative paths in complex environments under platform-specific constraints.

This paper examines the problem of active mapping using a robot, where the aim is to recover a continuous 2-D or 3-D field, e.g., of temperature, humidity, etc., using measurements collected by an on-board sensor. In similar setups, most existing strategies [1, 5, 8] incorrectly assume perfect pose information, which is an implicit requirement for building high-quality maps in initially unknown environments. Our motivation is to improve upon the robustness and accuracy of field reconstructions by allowing the robot to adaptively trade-off between gathering new information (exploration) and maintaining good localization (exploitation).

Despite recent efforts [2, 6, 9, 10], propagating both the localization and field map uncertainties into the planning framework in a principled manner remains an open challenge. A major issue arises due to the different ways in which the target field and robot pose are modeled. In particular, our work considers the task of mapping a field using a Gaussian Process (GP) with the robot pose represented as a multivariate Gaussian distribution. In this setup, uncertainty measures

[1]Smart Robotics Lab., Imperial College London. [2]Centre for Autonomous Systems, University of Technology Sydney. [3]Autonomous Systems Lab. (ASL), ETH Zürich. M. Popović was also with ASL during this work. m.popovic@imperial.ac.uk.

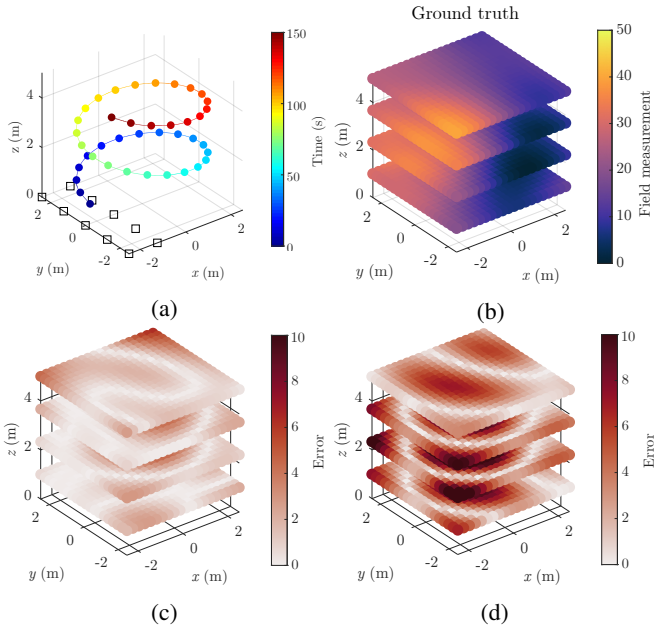

Fig. 1: Overview of our proposed active mapping strategy. (a) shows a spiral trajectory traveled by a robot. Squares indicate point landmarks on the ground used for localization. Spheres represent sites where measurements of the ground truth field map in (b) are taken. By accounting for pose estimation uncertainty, our framework yields an error map (c) with 2.47 times lower total error compared to a standard mapping approach (d).

in the field and robot pose, e.g., entropy-based criteria [4], are not directly comparable since they are obtained from their respective model variances which have different units and scales, and are therefore difficult to couple in a single utility function for multi-objective planning. Heuristic methods, e.g., based on a linear weighting of uncertainties [11], are commonly applied; however, they require careful manual tuning and are often scenario-specific.

To address this, we present an approach that accounts for the pose estimation uncertainty in two places. First, we consider the additional noise the pose uncertainty induces in the environmental field model; second, we include it as a shaping factor in the utility function that defines our informative planning task. Our mapping strategy uses GPs with uncertain inputs (UIs) [2] to propagate the pose uncertainty into the field model, as visualized in Fig. 1. During a mission, the map built online is used to plan informative trajectories in continuous space. We develop a new utility formulation for GP-based mapping scenarios that jointly considers the uncertainty of the robot and the field models. This enables us to capture the desired exploitation-exploration trade-off

in a mathematically sound manner, without relying on any manually-tuned, environment-dependent parameters. In summary, the contributions of this work are:

1) A utility function based on GP field models that tightly couples the objectives of robot localization and field mapping in active mapping problems.
2) An informative planning framework that accounts for the pose uncertainty in both mapping and the information objective for planning.
3) Evaluations of the approach in a 3-D graph Simultaneous Localization and Mapping (SLAM) setup and a proof of concept in a temperature mapping scenario.

We note that our framework can be used in any scalar field mapping scenario, e.g., spatial occupancy [4, 6, 12], signal strength [2, 8], aerial surveillance [5, 7], etc., and with any SLAM or localization-only algorithm supplying the pose uncertainty. Here, we focus on a general case where the spatial field data is decoupled from the localization routine.

II. RELATED WORK

Significant recent work has been done on autonomous data gathering in the context of robotics and related fields. The discussion here focuses on two main streams: (1) methods for probabilistic environmental mapping [2, 6, 9, 13], and (2) algorithms for informative planning [1, 4, 5, 12, 14, 15].

GPs are a popular non-parameteric Bayesian technique for modeling spatio-temporal phenomena [13]. They have been applied in various active sensing scenarios [1, 3, 5, 15] to gather data based on correlations and uncertainty in continuous maps. However, most of these works assume that the training data for prediction is inherently noise-free, which may lead to inaccuracies if measurements are incorporated at wrong locations and mislead predictive planning algorithms.

Propagating the input uncertainty through dense GP models is computationally challenging. Analytical [9] and heteroscedastic approximation methods [2, 3, 10] have been proposed, of which our work leverages the expected kernel technique of Jadidi et al. [2]. Specifically, we apply their approach to more complex 3-D planning scenarios, and integrate it with a new uncertainty-aware utility function. This coupling enables robust, tractable mapping under robot pose uncertainty for online sensing applications.

An active sensing task can be expressed in a Partially Observable Markov Decision Process (POMDP) [16] as one of decision-making under uncertainty. In practice, informative planning algorithms are typically used to render this problem computationally tractable with dense belief representations. We broadly distinguish between planning strategies operating in (a) discrete [4, 12] and (b) continuous space. This study focuses on the latter class of methods, which leverage incremental sampling [6, 8, 17, 18] or splines [1, 5] to offer greater scalability compared to discrete approaches. As in our prior work [5, 14, 15], we define smooth polynomial robot trajectories [19] and optimize them globally for an information objective in a finite-horizon manner.

Relatively limited research has been invested in active sensing scenarios where robot localization is uncertain. This setup has been tackled in the contexts of belief-space planning [6, 17, 18] and active SLAM [11, 12], where the aim is to maintain good localization while exploring an unknown environment. In contrast, and similarly to Papachristos et al. [6] and Costante et al. [17], our paper considers the map building and localization problems to be decoupled. The main distinction is that we aim to reconstruct a continuous field that is independent of the features used for localization.

An open challenge in this setup is formulating utility functions to trade off between robot localization and field mapping in a principled manner. Previous approaches have examined heuristic parameter tuning [12], e.g., using a linear weighting of the map and pose uncertainties [11], and multi-layer [6, 17] planning strategies. In contrast, we follow Carrillo et al. [4] in using the concept of Rényi's entropy to discount information gain based on predicted localization uncertainty. Thereby, our utility function shares the benefit of coupling the two objectives in a mathematically sound way, without manual tuning requirements. Our formulation is developed for a continuous mapping scenario based on a GP field model, instead of an occupancy grid [4]. Moreover, by mapping with UIs, we present a unified framework where the robot localization uncertainty is jointly accounted for in *both* mapping and planning to achieve robust data acquisition.

III. PROBLEM STATEMENT

The general active mapping problem is formulated as follows. We seek an optimal trajectory ψ^* in the space of all continuous trajectories Ψ to maximize an information-theoretic measure:

$$\begin{aligned} \psi^* = \underset{\psi \in \Psi}{\operatorname{argmax}} \, \mathrm{I}(\textsc{measure}(\psi)) \,, \\ \text{s.t. } \textsc{cost}(\psi) \leq B \,. \end{aligned} \tag{1}$$

The function $\textsc{measure}(\cdot)$ obtains a set of measurements along trajectory ψ, and $\textsc{cost}(\cdot)$ provides its associated cost, which cannot exceed a predefined budget B. The operator $\mathrm{I}(\cdot)$ defines the information objective quantifying the utility of new sensor measurements. In Sec. V-C, we propose a utility function for active mapping in GP-based scenarios that incorporates both the robot localization and field mapping objectives without manual parameter tuning requirements.

IV. MAPPING APPROACH

This section presents our mapping approach as the basis of our framework. We first describe our method for environmental field modeling using a GP, then present a strategy which folds the robot pose uncertainty into the map inference.

A. Gaussian Processes

We use a GP to model the spatial correlations of a field in a probabilistic and non-parametric manner [13]. The target field variable for mapping is assumed to be a continuous function: $f : \mathcal{E} \to \mathbb{R}$. A GP is characterized by a mean function $m(\mathbf{x}) \triangleq \mathbb{E}[f(\mathbf{x})]$ and covariance function $k(\mathbf{x}, \mathbf{x}') \triangleq \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}')]$ as $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$, where $\mathbb{E}[\cdot]$ is the expectation operator and $\mathbf{x}$ and $\mathbf{x}'$ are input vectors (spatial coordinates).

For a set of observed (training) locations in the fixed-size environment with associated target values, we apply the standard GP regression equations [13] to infer the mean $\boldsymbol{\mu}$ and covariance matrix $\mathbf{P}$ at a set of new query locations. The covariance function k encodes our assumptions about the field. In this work, we assume that the mapped environmental phenomena are smooth and isotropic; thus common choices for k are the squared exponential (SE) and Matérn [13].

B. Mapping Under Pose Uncertainty

To propagate the robot pose uncertainty into our mapping framework, we apply the expected kernel technique of Jadidi et al. [2] to planning problems in 3-D setups. The key idea lies in taking the expectation of k over UIs. Instead of X being a deterministic location as in Sec. IV-A, let X now be a random variable probabilistically distributed according to $p(x)$. The expected covariance function $\tilde{k}$ is computed as:

$$\tilde{k} = \mathbb{E}[k] = \int_X k p(x) dx \,. \tag{2}$$

Assuming a Gaussian distribution for the robot pose $\mathcal{N}(\mathbf{p}, \boldsymbol{\Sigma})$, we apply Gauss-Hermite quadrature to efficiently approximate the integral in Eq. (2) in three dimensions.

V. PLANNING APPROACH

Here we summarize the key steps of the informative planning algorithm, focusing on our uncertainty-aware utility function for GP-based active mapping scenarios as the main contribution of this paper. Interested readers are referred to our previous publications for more details [5, 14, 15].

A. Trajectories

A polynomial trajectory ψ is represented by N ordered control waypoints $\mathcal{C} = [\mathbf{c}_1, \dots, \mathbf{c}_N]$ connected using $N-1$ k-order spline segments. Given a reference velocity and acceleration, we optimize ψ for smooth minimum-snap dynamics [19], clamping $\mathbf{c}_1$ as the initial robot position. In Eq. (1), MEASURE$(\cdot)$ computes the spacing of measurement sites given a constant sensor frequency and robot speed.

B. Algorithm

We plan using a fixed-horizon approach, alternating replanning and execution until the elapsed time t exceeds the budget B. Our replanning strategy (Alg. 1) consists of two steps. First, an initial trajectory is obtained through a grid search (Lines 3-7) based on a coarse set of points $\mathcal{L}$ in the 3-D workspace. We conduct a sequential greedy search for N control waypoints; selecting the next-best point $\mathbf{c}^*$ (Line 4) by evaluating Eq. (1) over $\mathcal{L}$. As described in Sec. V-C, our new utility function (Eq. (5)) defines the objective $\mathrm{I}(\cdot)$ in Eq. (1) for planning under pose uncertainty $\boldsymbol{\Sigma}$. For each candidate goal, the evolution of $\boldsymbol{\Sigma}$ along a trajectory connecting it to the current pose is predicted (Line 5) according to Sec. V-D. The result is used to update the GP covariance matrix (Line 6) based on the new measurement locations. $\mathbf{c}^*$ is then added to the set of control points $\mathcal{C}$.

The second replanning step (Line 8) refines the coarse grid search output for $\mathcal{C}$ using Eq. (1). This is done by

Algorithm 1 REPLAN_PATH procedure

Input: Current covariance matrix of the GP field model $\mathbf{P}$, number of control waypoints N, grid points $\mathcal{L}$, initial position $\mathbf{c}_1$, robot pose $(\mathbf{p}, \boldsymbol{\Sigma})$

Output: Waypoints defining next polynomial plan $\mathcal{C}$

1: $\mathbf{P}' \leftarrow \mathbf{P}$; $(\mathbf{p}', \boldsymbol{\Sigma}') \leftarrow (\mathbf{p}, \boldsymbol{\Sigma})$ // Create local copies.
2: $\mathcal{C} \leftarrow \mathbf{c}_1$ // Initialize control points.
3: **while** $N \geq |\mathcal{C}|$ **do**
4: $\quad \mathbf{c}^* \leftarrow$ Select viewpoint in $\mathcal{L}$ using Eq. (1)
5: $\quad (\mathbf{p}', \boldsymbol{\Sigma}') \leftarrow$ PREDICT_MOTION$(\mathbf{p}', \boldsymbol{\Sigma}', \mathbf{c}^*)$
6: $\quad \mathbf{P}' \leftarrow$ PREDICT_MEASUREMENT$(\mathbf{P}', \mathbf{p}', \boldsymbol{\Sigma}')$
7: $\quad \mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}^*$
8: $\mathcal{C} \leftarrow$ CMAES$(\mathcal{C}, \mathbf{P}, \mathbf{p}, \boldsymbol{\Sigma})$ // Optimize control points using Eq. (1).

computing $\mathrm{I}(\cdot)$ for a sequence of measurements taken along a trajectory. For prediction, we apply the same principles as in Lines 5 and 6 above. A key benefit of our two-step approach is that the informed initialization speeds up the optimization convergence, making it suitable for finding solutions on computationally limited systems.

The optimization step is agnostic to the actual method considered; here, we employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [20], a global derivative-free optimizer based on evolutionary algorithms that has been used to solve high-dimensional, nonlinear, non-convex problems in continuous domains [1, 14, 15, 20]. Our choice of this routine is motivated by the nonlinearity of the objective in Eq. (1). Readers are referred to Hansen [20] for details.

C. Utility Definition

We introduce a new utility, or information gain, function $\mathrm{I}(\cdot)$ in Eq. (1) which directly relates the uncertainty of both the robot pose and the field map without the need for additional hand-tuned scaling parameters. Our utility function is derived from the Rényi entropy, which for a multivariate Gaussian distribution is given by [21]:

$$\mathrm{H}_\alpha(\mathbf{X}) = \frac{1}{2} \log \left(\left| \left(2\pi\alpha^{\frac{1}{\alpha-1}} \right) \mathbf{P} \right| \right), \tag{3}$$

where $\alpha \in [0, 1) \cup (1, \infty)$ is a free parameter. Note that the Rényi entropy converges to the Shannon entropy as $\alpha \to 1$.

To improve computational tractability, we approximate the determinant in Eq. (3) using the covariance matrix trace. Dropping the constant terms then gives:

$$\mathrm{H}_\alpha(\mathbf{X}) \propto\!\!\!\sim \hat{\mathrm{H}}_\alpha(\mathbf{X}) = \log \left(\mathrm{Tr}\left(\mathbf{P}\right) \alpha^{\frac{1}{\alpha-1}} \right), \tag{4}$$

where $\mathrm{Tr}(\cdot)$ is the matrix trace, and the covariance matrix $\mathbf{P}$ is obtained using Eq. (2) for mapping under uncertainty.

Our key insight is to relate the free parameter α to the predicted pose uncertainty $\boldsymbol{\Sigma}$ along a candidate trajectory such that the expected information gain reduces when localization uncertainty is high. Noting that H_α is strictly non-increasing in α, this intuition can be captured by setting α to $\alpha(\boldsymbol{\Sigma}) = 1 + \frac{1}{\mathrm{Tr}(\boldsymbol{\Sigma})}$, and computing information gain as:

$$\mathrm{I}_{\alpha(\boldsymbol{\Sigma})}(\mathbf{p}) = \hat{\mathrm{H}}(\mathbf{P}^-) - \hat{\mathrm{H}}_{\alpha(\boldsymbol{\Sigma})}(\mathbf{P}^+) \,. \tag{5}$$

The superscripts $-$ and $+$ denote the prior and posterior GP covariance matrix $\mathbf{P}$. Similar to our approximation in Eq. (4), the first term on the right side in Eq. (5) is computed using

the GP covariance trace. This approximates the Shannon entropy of the model before executing a candidate trajectory. The second term approximates the Rényi entropy of the model given the predicted measurements, which is shaped according to the associated pose uncertainties by $\alpha(\mathbf{\Sigma})$. Thus, Eq. (5) is similar to the mutual information, with the key difference that we explicitly discount information gain according to the predicted localization uncertainty.

Note that other functions can used to relate α to the localization uncertainty in $\alpha(\mathbf{\Sigma})$. Carrillo et al. [4] also suggest using the determinant or the maximum eigenvalue of $\mathbf{\Sigma}$ instead of the trace. Nevertheless, modulating the Rényi entropy provides a formal way in which to incorporate the localization uncertainty while maintaining a utility function that is primarily computed over the GP field uncertainty.

D. Uncertainty Prediction

A key requirement for predictive planning is propagating the robot localization uncertainty for a candidate action (Line 5 of Alg. 1). We consider this problem for active sensing in initially unknown and known environments.

1) Unknown Environments: We consider a solution assuming graph SLAM using odometry and point landmark observations [22]. Fig. 2 schematizes a 2-D example. To predict the localization uncertainty along a possible path (dashed line), the graph is simply extended from the current pose $\mathcal{N}(\mathbf{p}_1, \mathbf{\Sigma}_1)$. The trajectory is interpolated at a fixed frequency to add $K-1$ odometry constraints (hollow circles) in the extended graph, giving rise to the sequence $\{\mathcal{N}(\mathbf{p}_1, \mathbf{\Sigma}_1), \ldots, \mathcal{N}(\mathbf{p}_K, \mathbf{\Sigma}_K)\}$. For each consecutive node pair, we apply Gaussian noise with a variance proportional to the control input magnitude. This reflects the fact that longer steps are likely associated with higher actuation errors.

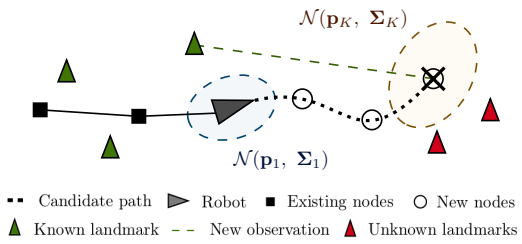

Fig. 2: Uncertainty prediction using graph SLAM with point landmarks. The final robot pose $\mathcal{N}(\mathbf{p}_K, \mathbf{\Sigma}_K)$ is estimated for planning.

To address potential loop closures, we simulate re-observations to the known landmarks (green triangles) maintained in the current graph given the predicted robot pose and sensor field of view (FoV). In unknown space, we assume no new landmark detections, so that the pose uncertainty grows due to control noise. The resulting graph is then solved using QR factorization to obtain $\mathbf{\Sigma}_K$ [23].

2) Known Environments: Similarly, uncertainty can be predicted assuming Monte Carlo Localization (MCL) in a known environment. This method uses a motion model and the expected sensor measurements to estimate the robot state by propagating the distribution of particles in a particle filter. It is applied in Sec. VI-C.

VI. EXPERIMENTAL RESULTS

A. Comparison of Planning Methods

First, the aim is to evaluate our new utility function by comparing it against existing strategies. To focus on examining our utility function and planning performance in particular, all methods in this sub-section use the GP-based approach with UIs for mapping under pose uncertainty (Sec. IV-B). The experiments are executed in MATLAB running on an Intel 1.8 GHz computer with 16 GB of RAM.

Our setup considers 5 different $5\,\text{m}\times 5\,\text{m}\times 4\,\text{m}$ Gaussian Random Field environments for mapping, assumed to be initially unknown. We use a $0.25\,\text{m}\times 0.25\,\text{m}\times 1\,\text{m}$ resolution grid and apply the isotropic SE kernel with hyperparameters trained by minimizing log marginal likelihood in each environment [13]. For mapping with UIs, the modified kernel in Eq. (2) is estimated using Gauss-Hermite quadrature with 5 points. During a mission, field measurements are taken at 0.25 Hz using a point-based sensor centered on a 3 degrees of freedom robot moving as a point mass along the (x, y, z) axes, and are added to our GP model as input training points to achieve uncertainty reduction as exploration takes place.

To emulate an aerial robot setup, 10 visual 3-D point landmarks are placed on the ground 1 m below the target field and distributed over one side of the space, as shown in Figs. 1a and 5. For SLAM, the robot is equipped with a downward-facing camera with $(47.9°, 36.9°)$ FoV. Landmark detection is based on a pinhole projection camera model with standard deviations of 1.0 px and 0.1 m for errors in pixel and depth. Our framework uses graph SLAM [22] with the approach in Sec. V-D.1 to predict localization uncertainty. We sample trajectories at 0.5 Hz to simulate control actions using an odometry motion model, applying a coefficient of 0.01 in all three dimensions to scale the noise variance.

Our two-step CMA-ES-based replanning framework (Sec. V-B) with the Rényi-based utility function (Sec. V-C) is compared against itself using different objectives: (a) field map uncertainty reduction only (Eq. (5) using the Shannon entropy in the posterior); (b) field map uncertainty reduction over time (rate), as in our previous works [5, 14, 15]; and (c) a linear composite weighting the map and pose uncertainties equally based on their upper bounds, tuned according to Bourgault et al. [11]. As benchmarks, we also study: the rapidly exploring information gathering tree (RIG-tree) [8], a state-of-the-art sampling-based informative planner, using objectives of (d) uncertainty reduction only and (e) our Rényi-based utility, as well as (f) random planning. We evaluate against the RIG-tree to assess our two-step replanning routine and, in (e), show the applicability of our generic concepts in uncertainty-aware active mapping with different planners. A 150 s budget B is fixed for all methods. We evaluate mapping uncertainty with the GP covariance trace $\text{Tr}(\mathbf{P})$ and accuracy with Root Mean Squared Error (RMSE) compared to the ground truth map. Similarly, for planning, our measures are the robot covariance trace $\text{Tr}(\mathbf{\Sigma})$ and the pose RMSE compared to the ground truth trajectory. Lower values using all metrics signify better performance.

The robot starts at (2 m, 2 m, 1 m) with no prior localization nor field map information. For trajectory optimization, the reference velocity and acceleration are $1.5\,\mathrm{m/s}$ and $3\,\mathrm{m/s^2}$ using polynomials of order $k = 12$. In our planner, we define polynomials with $N = 4$ waypoints and use a uniformly spaced 27-point lattice for the 3-D grid search. In the RIG-tree, we associate waypoints with vertices and form polynomials by tracing their parents. The finite-horizon replanning procedure follows the approach of Popović et al. [5] and a branch step-size of 5 m is set for best performance based on a search over a discrete value range. We set 60 sampling iterations to obtain the same ~ 76 s replanning time as required by our routine. In the random planner, random destinations are sampled in the workspace and a trajectory is generated by connecting them sequentially to the current position. We consider 4 waypoints per plan to ensure that trajectory lengths are fairly comparable to our method.

We conduct 50 trials in each environment. Fig. 3 shows how the metrics evolve for each method. As expected, in Fig. 3b, informed strategies perform better than the random benchmark (blue). In Fig. 3a, the uncertainty-based (black, green) and weighted (cyan) utility functions yield map uncertainty and error reduction rates similar to our Rényi-based objective (red). However, our method significantly improves upon the pose estimation; effectively trading off between gathering information and maintaining good localization. This cannot be done using manually-tuned parameters (cyan) as the variability of the pose uncertainty is much lower compared to that of the map. In terms of localization, the uncertainty-only function (black) is the worst as it does not exploit any knowledge of the trajectory dynamics.

In Fig. 3b, the RIG-tree using our Rényi-based (purple) objective provides better localization compared to planning for map uncertainty reduction only (yellow). This further validates the behaviour of our utility function. Interestingly, the RIG-tree scores comparatively well on the localization metrics while the map improves at a limited rate. This likely relates to the step-size parameter the algorithm requires. By traveling in shorter steps, the robot re-observes landmarks more frequently at the cost of restricted exploration.

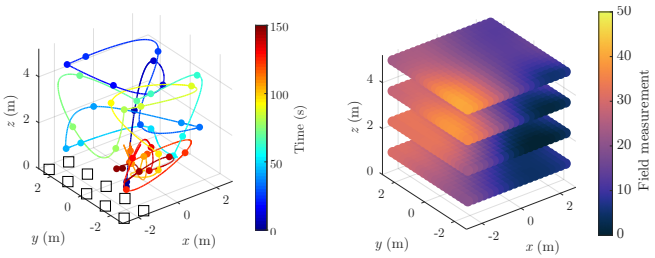

Fig. 5: Example result using our CMA-ES-based framework and Rényi-based utility. The left plot shows the traveled trajectory. Squares indicate point landmarks on the ground used for localization. Spheres represent sites where measurements are taken to produce the field map on the right. Our planner effectively balances between exploring new areas and keeping landmarks in view.

Fig. 5 exemplifies a result using our proposed approach. The left plot confirms that the robot explores the space while re-visiting landmarks to stay well-localized. With a total RMSE of 1.11, our final map (right) is 1.86 times more accurate than the one output by the naïve spiral path in Fig. 1.

B. Evaluation of Field Mapping Under Uncertainty

Next, the aim is to assess the benefits of mapping under pose uncertainty by evaluating the effects of incorporating UIs in the GP field model. We consider the same simulation setup as above using our CMA-ES-based replanning framework and Rényi-based utility function. We conduct 50 trials in each of the 5 environments (a) with and (b) without applying the modified kernel in Sec. IV-B. The benchmark (b) corresponds to standard GP mapping without UIs.

Fig. 6 depicts our results. Note that we omit the map uncertainty metric as the variance scales using the two methods

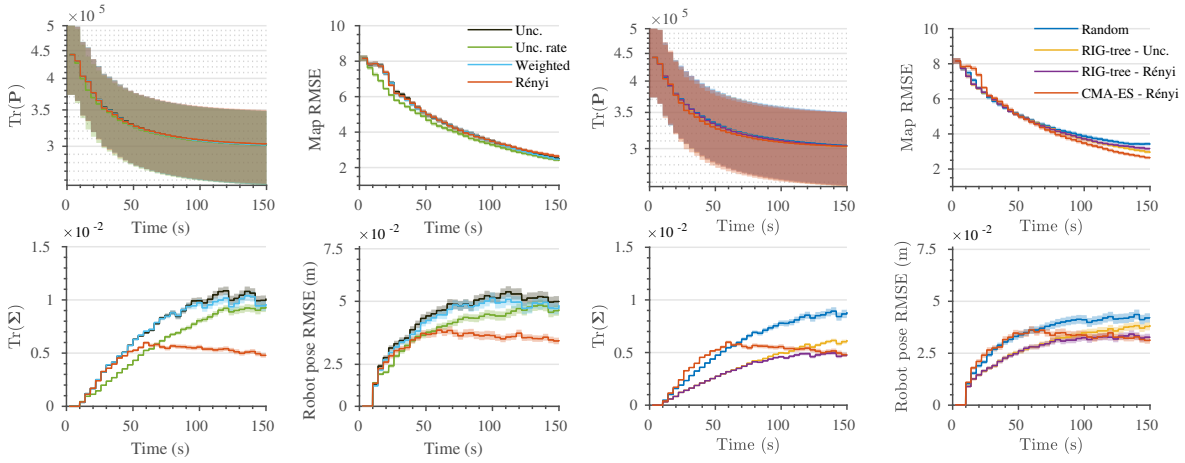

Fig. 3: Comparison of our two-step CMA-ES-based replanning framework using our Rényi-based utility against (a) our framework using different information objectives and (b) different planning benchmarks. All methods are given a 150 s budget and use GP-based field mapping with UIs. Solid lines represent means over 250 trials. Shaded regions show 95% confidence bounds. By considering both the robot pose and field map uncertainties, our utility function more quickly achieves higher-quality mapping (top) with improved localization (bottom). Note that the $\mathrm{Tr}(\mathbf{P})$ axis is logarithmic and that it has high variance due to different field environments studied in the trials.

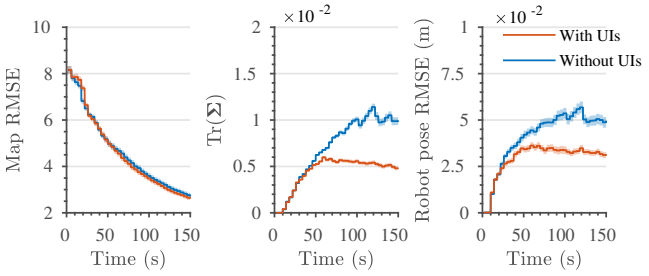

Fig. 6: Comparison of field mapping with and without uncertain inputs (UIs) using our CMA-ES-based replanning framework and Rényi-based utility for a 150 s budget. Solid lines represent means over 250 trials. Shaded regions show 95% confidence bounds. By accounting for pose uncertainty, our approach achieves more conservative mapping (middle, right) with higher accuracy (left).

are not comparable. The plots confirm that our approach with UIs (red) presents more conservative exploratory behavior than the benchmark (blue) while also yielding more accurate field reconstructions. This is because the modified kernel can handle localization errors to build more consistent maps for planning. Fig. 1 shows a visualization supporting this result.

C. Proof of Concept

We show our framework mapping in real-time on a TurtleBot3 Waffle with an Intel Joule 570x running Ubuntu Linux 16.04 and the Robot Operating System. The experiments transpire in a known indoor environment using Adaptive Monte Carlo Localization (AMCL) receiving data from a LDS-01 laser scanner. As shown in Fig. 7, a temperature distribution in an empty $2.8\,\mathrm{m} \times 2.8\,\mathrm{m}$ area is generated using a 2400 W radiant heater placed at one corner. Measurements are taken using a LM35 linear temperature sensor with a sensitivity of $10\,\mathrm{mV/°C}$.

A $0.4\,\mathrm{m}$ resolution grid is set for mapping with UIs. To train the GP, we follow the method in Sec. VI-B using manually gathered data to obtain the hyperparameters for the SE kernel. The integral in Eq. (2) is estimated using 5 Gauss-Hermite points. Uncertainty prediction is based on localization in a known environment using AMCL (Sec. V-D). We sub-sample each candidate plan at 2 Hz and estimate pose using a differential drive odometry model and laser scans simulated in the known occupancy map. The control noise variance is $0.2\,\mathrm{m}^2$ for both rotation and translation.

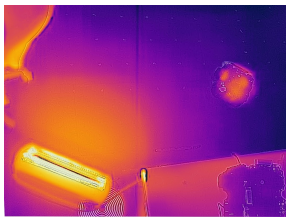

Fig 7: Thermal imagery of our experimental setup showing the robot and radiator from an aerial view. Yellower shades correspond to heated areas mapped using the onboard temperature sensor.

The aim is to show our framework mapping a realistic continuous field. The initial measurement point is $(0.2\,\mathrm{m}, 0.2\,\mathrm{m})$ within the corner opposite the radiator. We allocate a budget B of 600 s. Following the differential drive model, each plan is piecewise linear as defined by $N = 3$ waypoints with a constant velocity of $0.26\,\mathrm{m/s}$ and temperature measurements sampled at $0.25\,\mathrm{Hz}$. The planning objective is our Rényi-based utility function (Eq. (5)). Note that field map updates are triggered upon allowing the sensor readings to stabilize.

Fig. 4 summarizes our experiments. As expected, the map becomes more complete and uncertainty decreases as the yellower heated region is discovered. Note that the values of the measurements are higher than that at which the field is initialized (23.64°). This is due to the effects of heating and diffusion over time, after the training data was collected.

VII. CONCLUSIONS AND FUTURE WORK

This work introduced an informative planning framework for active mapping that accounts for the robot pose uncertainty in both the mapping and planning stages. Our method uses GPs with UIs to propagate pose uncertainty into the model of an environmental field. For planning, we proposed a new utility function that tightly couples the uncertainties in the robot pose and field map by applying the concept of Rényi's entropy in GP-based mapping scenarios. Our formulation addresses the exploration-exploitation trade-off in a principled way, without manually-tuned parameters.

Our framework was evaluated extensively in simulation. We showed that it achieves more conservative exploratory behavior compared to different strategies while producing more accurate maps. Experimental validation was performed through a proof of concept deployment. Future work will examine temporal field models and efficiency improvements towards more complex environments. Other interesting research directions involve refining the uncertainty prediction method and its relationship to α for more reliable planning.

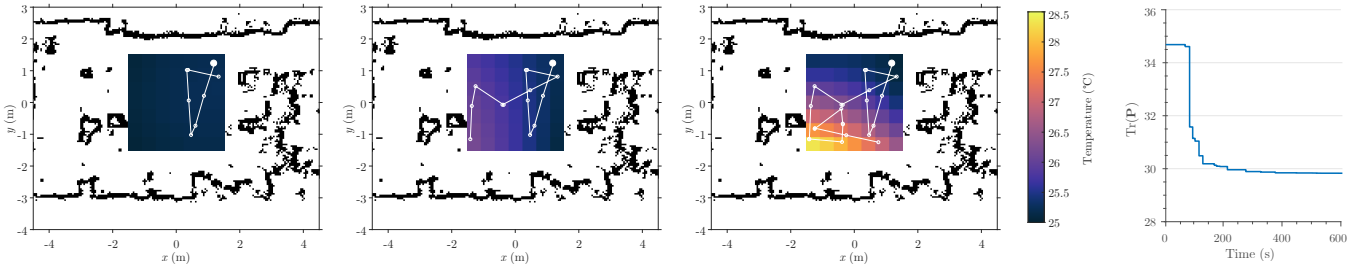

Fig. 4: Results of using our framework to map the temperature distribution (Fig. 7) in a 600 s mission. The three left plots depict the trajectories (white lines) and temperature field maps (colored gradients) at $t = 100$ s, 350 s, and 600 s. White circles represent measurement sites, with the large solid one indicating the initial robot position. The sequence shows that our planner quickly explores the area; detecting the corner heated by the radiator (bottom-left). The right curve validates map uncertainty reducing over time. Planning time is included.

REFERENCES

[1] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive Continuous-Space Informative Path Planning for Online Environmental Monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017.

[2] M. G. Jadidi, J. V. Miro, and G. Dissanayake, "Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring," *The International Journal of Robotics Research*, vol. 38, no. 6, pp. 658–685, 2019.

[3] R. Oliveira, L. Ott, V. Guizilini, F. Ramos, and R. O. Sep, "Bayesian Optimisation for Safe Navigation under Localisation Uncertainty," in *International Symposium of Robotics Research*. Puerto Varas: Springer, 2017.

[4] H. Carrillo, P. Dames, V. Kumar, and J. A. Castellanos, "Autonomous robotic exploration using a utility function based on Rényi's general theory of entropy," *Autonomous Robots*, vol. 42, no. 2, pp. 235–256, 2018.

[5] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for UAV-based terrain monitoring," *Autonomous Robots*, 2020.

[6] C. Papachristos, S. Khattak, and K. Alexis, "Uncertainty-aware Receding Horizon Exploration and Mapping using Aerial Robots," in *IEEE International Conference on Robotics and Automation*. Singapore: IEEE, 2017, pp. 4568–4575.

[7] S. Manfreda, M. F. McCabe, P. E. Miller, R. Lucas, V. P. Madrigal, G. Mallinis, E. B. Dor, D. Helman, L. Estes, G. Ciraolo, J. Müllerová, F. Tauro, M. I. de Lima, J. L. de Lima, A. Maltese, F. Frances, K. Caylor, M. Kohv, M. Perks, G. Ruiz-Pérez, Z. Su, G. Vico, and B. Toth, "On the Use of Unmanned Aerial Systems for Environmental Monitoring," *Remote Sensing*, vol. 10, no. 4, 2018.

[8] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.

[9] A. Girard, "Approximate methods for propagation of uncertainty with Gaussian process models," Ph.D. dissertation, University of Glasgow, 2004.

[10] A. Mchutchon and C. E. Rasmussen, "Gaussian Process Training with Input Noise," *Advances in Neural Information Processing Systems*, pp. 1341–1349, 2011.

[11] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, "Information Based Adaptive Robotic Exploration," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Lauseanne: IEEE, 2002, pp. 540–545.

[12] R. Valencia, J. V. Miró, G. Dissanayake, and J. Andrade-Cetto, "Active Pose SLAM," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vilamoura, 2012, pp. 1885–1891.

[13] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.

[14] M. Popović, G. Hitz, J. Nieto, I. Sa, R. Siegwart, and E. Galceran, "Online Informative Path Planning for Active Classification Using UAVs," in *IEEE International Conference on Robotics and Automation*. Singapore: IEEE, 2017.

[15] M. Popović, T. Vidal-Calleja, G. Hitz, I. Sa, R. Y. Siegwart, and J. Nieto, "Multiresolution Mapping and Informative Path Planning for UAV-based Terrain Monitoring," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Vancouver: IEEE, 2017.

[16] L. Kaelbling, M. Littman, and A. Cassandra, "Planning and Acting in Partially Observable Stochastic Domains," *Artificial Intelligence*, vol. 101, no. 1-2, pp. 99–134, 1998.

[17] G. Costante, J. Delmerico, M. Werlberger, and P. Valigi, "Exploiting Photometric Information for Planning under Uncertainty," in *Robotics Research*. Springer, 2017, pp. 107–124.

[18] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *IEEE International Conference on Robotics and Automation*. Shanghai: IEEE, 2011, pp. 723–730.

[19] C. Richter, A. Bry, and N. Roy, "Polynomial Trajectory Planning for Aggressive Quadrotor Flight in Dense Indoor Environments," in *International Symposium of Robotics Research*. Singapore: Springer, 2013.

[20] N. Hansen, "The CMA evolution strategy: A comparing review," *Studies in Fuzziness and Soft Computing*, vol. 192, no. 2006, pp. 75–102, 2006.

[21] L. Golshani and E. Pasha, "Rényi entropy rate for Gaussian processes," *Information Sciences*, vol. 180, no. 8, pp. 1486–1491, 2010.

[22] J. Solà, "Course on SLAM," Barcelona, 2017.

[23] M. Kaess and F. Dellaert, "Covariance recovery from a square root information matrix for data association," *Robotics and Autonomous Systems*, vol. 57, no. 12, pp. 1198–1210, 2009.