

Informative Path Planning for Active Field Mapping under Localization Uncertainty

Marija Popović, Teresa Vidal-Calleja, Jen Jen Chung, Juan Nieto, Roland Siegwart

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 any manually tuned parameters. Extensive simulations show that our approach outperforms existing strategies, with reductions in mean pose uncertainty and map error. We also 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

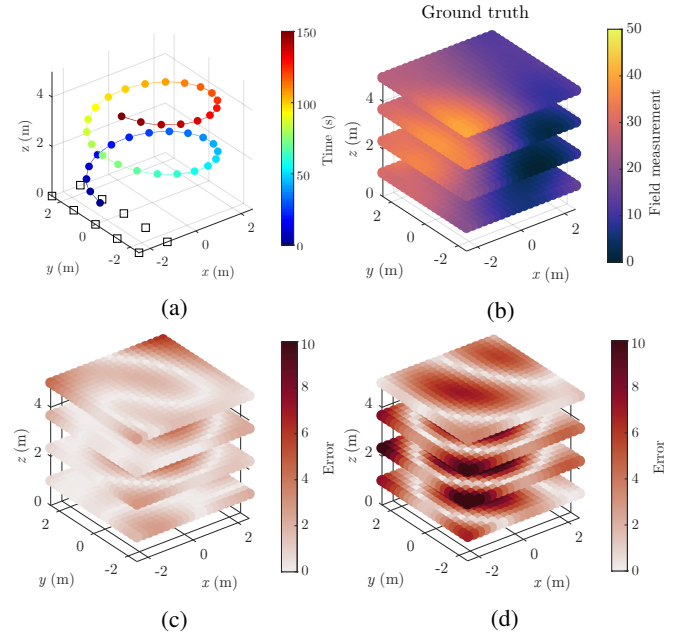


Fig. 1: Overview of our proposed active mapping strategy. (a) shows a spiral trajectory traveled by a robot. The squares indicate point landmarks on the ground used for localization. The spheres represent sites where measurements of the ground truth field map in (b) are taken. By accounting for pose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

Gaussian Process (GP) with the robot pose represented as a multivariate Gaussian distribution. In this setup, measures of uncertainty 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 Figure 1. During a mission, the map built online is used to plan informative trajectories in continuous space by optimizing initial solutions obtained by

M. Popović, J. J. Chung, J. Nieto, and R. Siegwart are with the Autonomous Systems Lab., ETH Zürich, Zürich, Switzerland. T. Vidal-Calleja is with the Centre for Autonomous Systems at the Faculty of Engineering and IT, University of Technology Sydney, Australia. Corresponding author: mpopovic@ethz.ch.

a coarse grid search. 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. Specifically, in this work, we focus on a general case where the spatial field information being collected is decoupled from the localization routine. More generally, the applicability of our utility function extends to other areas of robotics, such as reinforcement learning [13].

II. RELATED WORK

Significant recent work has been done on autonomous data gathering strategies in the context of robotics and related fields. The discussion in this section focuses on two main research streams: (1) methods for probabilistic environmental mapping [2, 6, 9, 14], and (2) algorithms for informative planning [1, 4, 5, 12, 15, 16].

GPs are a popular non-parameteric Bayesian technique for modeling spatio-temporal phenomena [14]. They have been applied in various active sensing scenarios [1, 3, 5, 16] 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 a computationally challenging task. To address this, analytical [9] and heteroscedastic approximation methods [2, 3, 10] have been proposed. Our work leverages the expected kernel technique of Jadidi et al. [2] by integrating over UIs with deterministic query points. Specifically, we apply their approach to more complex planning problems in 3-D 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) [17] 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, 18, 19] or splines [1, 5] to offer greater scalability compared to discrete approaches. As in our prior work [5, 15, 16], we define smooth polynomial robot trajectories [20] 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, 18, 19] and active SLAM [11, 12], where the aim is to maintain good localization as an unknown environment is explored. In contrast, and similarly to Papachristos et al. [6] and Costante et al. [18], our paper considers a setting where the map building and robot localization problems are decoupled. An important 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. To address this, previous approaches have examined heuristic parameter tuning [12], e.g., using a weighted linear combination of the map and pose uncertainties [11], and multi-layer [6, 18] planning strategies. In contrast to these methods, 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. The core difference is that our formulation is developed for a continuous mapping scenario based on a GP field model, instead of an occupancy grid. 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 more 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:

$$\begin{aligned} \psi^* &= \underset{\psi \in \Psi}{\operatorname{argmax}} I(\text{MEASURE}(\psi)), \\ \text{s.t. } \text{COST}(\psi) &\leq B. \end{aligned} \quad (1)$$

The function $\text{MEASURE}(\cdot)$ obtains a finite set of measurements along trajectory ψ , and $\text{COST}(\cdot)$ provides its associated cost, which cannot exceed a predefined budget B . The operator $I(\cdot)$ defines the information objective quantifying the utility of the acquired sensor measurements. In [Section 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 any 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 spatial correlations of a field in a probabilistic and non-parametric manner [14]. The target field variable for mapping is assumed to be a continuous function: $f : \mathcal{E} \rightarrow \mathbb{R}$. Essentially, a GP is a generalization of the Gaussian probability distribution, where the stochastic process governs the properties of functions, as opposed to vectors or scalars in the case of the probability distribution. A GP is fully characterized by the 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 composed of spatial coordinates in the environment.

Let $\mathbf{X} \subset \mathcal{E}$ be a set of n observed robot locations \mathbf{x}_i in the fixed-size environment with associated target field values \mathbf{y} . For a new set of n^* query locations $\mathbf{X}_* \subset \mathcal{E}$ we can infer the field Gaussian distribution by conditioning in the observed values $\mathbf{f}_* | \mathbf{X}, \mathbf{y}, \mathbf{X}_* \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{P})$ as follows [14]:

$$\boldsymbol{\mu} = m(\mathbf{X}_*) + K(\mathbf{X}_*, \mathbf{X}) [K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}]^{-1} \times (\mathbf{y} - m(\mathbf{X})), \quad (2)$$

$$\mathbf{P} = K(\mathbf{X}_*, \mathbf{X}_*) - K(\mathbf{X}_*, \mathbf{X}) [K(\mathbf{X}, \mathbf{X}) + \sigma_n^2 \mathbf{I}]^{-1} \times K(\mathbf{X}, \mathbf{X}_*)^\top, \quad (3)$$

where σ_n^2 is a hyperparameter representing the observation noise variance, $K(\mathbf{X}_*, \mathbf{X})$ denotes the cross-correlation terms between the query and observed points, and $K(\mathbf{X}, \mathbf{X})$ and $K(\mathbf{X}_*, \mathbf{X}_*)$ are the joint covariance matrices for the observed and query points, respectively. Note that $K(\cdot, \cdot)$ corresponds to the covariance function $k(\cdot, \cdot)$ for only one element.

Note that k encodes our assumptions about the field we are interested in modeling. In this work, we assume that the mapped environmental phenomena is smooth and isotropic; thus common choices of covariance functions are the squared exponential (SE) and Matérn family [14].

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 Section Section IV-A, let X now be a random variable distributed according to a probability distribution $p(x)$. The expected covariance function \tilde{k} can be computed as:

$$\tilde{k} = \mathbb{E}[k] = \int_{\mathcal{X}} kp(x)dx. \quad (4)$$

Assuming a Gaussian distribution for the robot position $\mathcal{N}(\mathbf{p}, \boldsymbol{\Sigma})$, we apply Gauss-Hermite quadrature to efficiently

approximate the integral in Equation (4) in three dimensions. Specifically, in this case, through a change of variable such that $\mathbf{L}\mathbf{L}^\top = \boldsymbol{\Sigma}$ and $\mathbf{u} = \mathbf{L}^{-1}(\mathbf{x} - \mathbf{p})$, where \mathbf{L} is a lower triangular matrix computed via the Cholesky decomposition:

$$\tilde{k} = \frac{1}{(2\pi)^{\frac{3}{2}}} \sum_{i_1=1}^M \sum_{i_2=1}^M \sum_{i_3=1}^M \bar{w} k_{i_1:i_3}, \quad (5)$$

where M is the number of sample points used in the approximation, $\bar{w} \triangleq \prod_{j=1}^3 w_{i_j}$, u_{i_j} corresponds to the roots of the Hermite polynomial H_n , $\mathbf{u}_{i_1:i_3} \triangleq [u_{i_1}, u_{i_2}, u_{i_3}]^\top$, and $k_{i_1:i_3}$ is the covariance function evaluated at $\mathbf{x}_{i_1:i_3} = \mathbf{L}\mathbf{u}_{i_1:i_3} + \mathbf{p}$.

V. PLANNING APPROACH

This section overviews our informative planning scheme, which generates fixed-horizon plans through a combination of a 3-D grid search and evolutionary optimization. We summarize the key steps of the algorithm, focusing on our uncertainty-aware utility function for GP-based active mapping scenarios as the main contribution of this paper. For further details concerning the planning approach itself, the reader is referred to our previous publications [5, 15, 16].

A. Trajectories

A polynomial trajectory ψ is represented by N ordered control waypoints to visit $\mathcal{C} = [\mathbf{c}_1, \dots, \mathbf{c}_N]$ connected using $N - 1$ k -order spline segments. Given a reference velocity and acceleration, we optimize the trajectory for smooth minimum-snap dynamics [20], clamping \mathbf{c}_1 as the initial robot position. In Equation (1), MEASURE(\cdot) is defined by computing the spacing of measurement sites given a constant sensor frequency and the traveling speed of the robot.

B. Algorithm

We plan using a fixed-horizon approach, alternating between replanning and execution until the elapsed time t exceeds the budget B . Our replanning strategy (Algorithm 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 robot workspace. In this step, we conduct a sequential greedy search for N control waypoints in \mathcal{C} ; selecting the next-best point \mathbf{c}^* (Line 4) by evaluating Equation (1) over \mathcal{L} . As described in Section V-C, our new utility function (Equation (9)) defines the objective $I(\cdot)$ in Equation (1) for planning under pose uncertainty. During this procedure, for each candidate goal, the evolution of the robot pose uncertainty $\boldsymbol{\Sigma}$ along a trajectory connecting it to the current pose is predicted using SLAM (Line 5), as detailed in Section V-D. The result is used to update the covariance matrix of the GP field model \mathbf{P} (Line 6) via Equation (3) based on the locations at which new measurements are taken. It is then added to the points \mathcal{C} representing the initial trajectory solution (Line 7).

The second replanning step (Line 8) refines the coarse grid search output for \mathcal{C} using Equation (1). This is done by computing $I(\cdot)$ for a sequence of measurements taken along the corresponding trajectory. For prediction, we apply the same principles as described above in Lines 5 and

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}, Σ)

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

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

6. An important benefit of our two-step approach is that the informed initialization procedure exploiting the grid-based solution effectively speeds up the convergence of the optimizer, making it suitable for quickly finding solutions on computationally limited systems.

Note that the optimization step is agnostic to the actual method considered; specifically, in this work, we employ the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) [21]. The CMA-ES is a generic global derivative-free optimizer based on evolutionary algorithms that has been successfully applied to solve high-dimensional, nonlinear, non-convex problems in a continuous domain [1, 15, 16, 21]. Essentially, it operates by iteratively evolving candidate solutions drawn from a search distribution whose parameters are adapted over time. Our choice of this routine is motivated by the nonlinearity of the objective space in Equation (1) as well as by the above mentioned results. For further details, the reader is referred to the in-depth review of Hansen [21].

C. Utility Definition

We introduce a new utility, or information gain, function $I(\cdot)$ in Equation (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 [22]:

$$H_\alpha(\mathbf{X}) = \frac{1}{2} \log \left(\left| \left(2\pi\alpha^{\frac{1}{\alpha-1}} \right) \mathbf{P} \right| \right), \quad (6)$$

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

To improve computational tractability, we approximate the determinant in Equation (6) using the trace of the covariance matrix [1, 16, 18]. Dropping the constant terms then gives:

$$H_\alpha(\mathbf{X}) \propto \hat{H}_\alpha(\mathbf{X}) = \log \left(\text{Tr}(\mathbf{P}) \alpha^{\frac{1}{\alpha-1}} \right), \quad (7)$$

where $\text{Tr}(\cdot)$ denotes the matrix trace, and the covariance matrix \mathbf{P} is obtained using the expected kernel in Equation (4) for mapping under uncertainty.

Our key insight is to relate the free parameter α to the predicted pose uncertainty Σ 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:

$$\alpha(\Sigma) = 1 + \frac{1}{\text{Tr}(\Sigma)}, \quad (8)$$

and computing information gain as:

$$I_{\alpha(\Sigma)}(\mathbf{p}) = \hat{H}(\mathbf{P}^-) - \hat{H}_{\alpha(\Sigma)}(\mathbf{P}^+), \quad (9)$$

where the superscripts $-$ and $+$ denote the prior and posterior covariance matrix \mathbf{P} of the GP model, respectively. Similar to our approximation in Equation (7), the first term on the right hand side in Equation (9) is computed using the trace of the GP covariance. This approximates the Shannon entropy of the model before executing the candidate trajectory. The second term approximates the Rényi entropy of the model given the predicted observations, which is shaped according to the associated pose uncertainties by Equation (8). Thus, Equation (9) bears similarities to the mutual information, with the key difference that we are able to explicitly discount the information gain according to the predicted localization uncertainty along the candidate path.

It is worth noting that other functions may be used to relate α to the localization uncertainty. For example, Carrillo et al. [4] also suggest using the determinant or the maximum eigenvalue of Σ in place of the trace in Equation (8). 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 map uncertainty.

D. Uncertainty Prediction

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

1) *Unknown Environments*: In unknown environments, we consider a solution to this task assuming graph SLAM using odometry and point landmark observations as constraints [23]. Figure 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 $\mathcal{N}(\mathbf{p}_1, \Sigma_1)$. 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 $\{\mathcal{N}(\mathbf{p}_1, \Sigma_1), \dots, \mathcal{N}(\mathbf{p}_K, \Sigma_K)\}$. For each consecutive node pair, we apply control noise drawn from a Gaussian distribution whose variance is proportional to the control input magnitude. This reflects the fact that longer motion steps are likely to be associated with higher actuation errors.

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 that no new landmarks will be detected, and thus, the robot pose uncertainty will grow as a result of the control noise. The resulting graph is then solved using QR factorization to




Fig. 2: Uncertainty prediction using graph SLAM with point landmarks. The dashed line is a candidate path to ‘X’. Example uncertainty ellipses at initial and final positions are shown. To predict the robot pose evolution, new nodes are interpolated along the path to extend the current graph. The extended graph is optimized to estimate the final pose $\mathcal{N}(\mathbf{p}_K, \Sigma_K)$ for planning.

obtain Σ_K . Future work will consider extending these ideas to other SLAM systems.

2) *Known Environments*: Similarly, in a known environment, the uncertainty can be predicted assuming a Monte Carlo Localization (MCL) approach. This method uses a motion model and the expected sensor measurements obtained to estimate the robot state by propagating the distribution of particles in a particle filter. In Section VI-C, this method is applied to localize a robot with a laser scanner in an indoor area.

Putting together our ideas in mapping and planning in a single framework, the GP field model with UIs accounts for the pose uncertainty in the received observations, while the proposed utility function accounts for the pose uncertainty in potential future measurements. This coupling allows us to achieve robust performance in GP-based mapping scenarios.

VI. EXPERIMENTAL RESULTS

In this section, our approach is evaluated in simulation by comparing it to different strategies for planning and mapping. We then show proof of concept by using it to map temperature using a ground robot in an indoor environment.

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 (Section IV-B). The experiments are executed in MATLAB running on an Intel 1.8 GHz computer with 16 GB of RAM.

Our experimental 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 [14]. For mapping with UIs, the modified kernel in Equation (4) 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. These new measurements are added to our GP model as observed 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 are distributed over one side of the space, as shown in Figures 1a and 4a. For SLAM, the robot is equipped with a downward-facing camera with $(47.9^\circ, 36.9^\circ)$ FoV. Landmark detection is done based on a pinhole projection camera model with standard deviations of 1.0 px and 0.1 m for measurement errors in pixel and depth. Our framework uses the graph SLAM algorithm [23] implemented in the SLAM Toolbox¹, with the approach described in Section V-D.1 for predicting the robot localization uncertainty in unknown environments. 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 co-ordinate dimensions to scale the control noise variance. For planning, covariance matrices from the graph are extracted through the computations developed by Kaess and Dellaert [24].

Our two-step CMA-ES-based replanning framework (Section V-B) with the Rényi-based utility function (Section V-C) is compared against itself using different objectives: (a) field map uncertainty reduction only (Equation (9) using standard Shannon’s entropy in the map posterior); (b) field map uncertainty reduction over time (rate), as in our previous works [5, 15, 16]; and (c) a linear composite of the map and pose uncertainties weighting both objectives equally based on their upper bounds, tuned according to the approach of Bourgault et al. [11]. As planning benchmarks, we also study: the rapidly exploring information gathering tree (RIG-tree) using objectives of (d) uncertainty reduction only and (e) our Rényi-based utility, as well as (f) random waypoint selection. The former is a state-of-the-art sampling-based planning algorithm introduced by Hollinger and Sukhatme [8] that incrementally extends a geometric tree to find the best information gathering trajectory. We consider this benchmark to assess our two-step replanning approach and, in (f), show the applicability of our generic concepts in uncertainty-aware active sensing with different planners. A 150 s budget B is specified for all methods. To evaluate mapping performance, we quantify uncertainty with the GP covariance trace $\text{Tr}(\mathbf{P})$ and accuracy with Root Mean Squared Error (RMSE) with respect to the ground truth map. Similarly, for planning, our measures are the robot covariance trace $\text{Tr}(\Sigma)$ and the pose RMSE with respect to the ground truth trajectory. Intuitively, lower values using all metrics signify better performance.

In all environments, the starting robot position is (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 m/s and 3 m/s² 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 RIG-tree, we associate control waypoints

¹github.com/joansola/slamtb

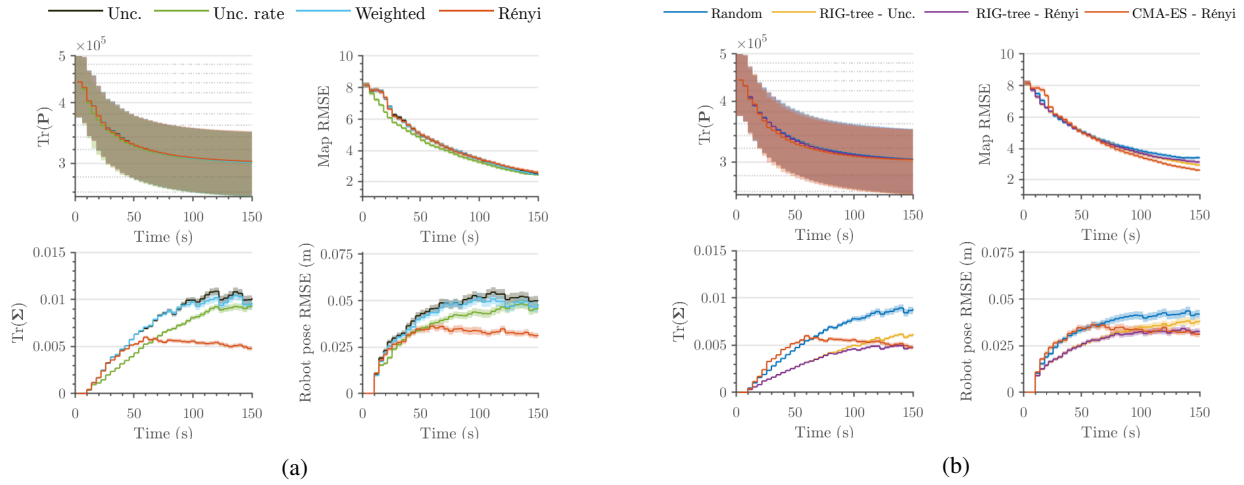


Fig. 3: Comparison of our two-step CMA-ES-based replanning framework using our Rényi-based utility function against (a) our framework using different information objectives and (b) different planning benchmarks. All methods are given a fixed budget of 150 s and use GP-based field mapping with UIs. The solid lines represent means over 250 trials. The 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 row) with improved localization (bottom row). Note that $\text{Tr}(\mathbf{P})$ is on a logarithmic scale, and that its variance is large due to differences between the field mapping environments considered in the trials.

with vertices, and form polynomials by tracing the parents of leaf vertices to the root. The finite-horizon replanning procedure follows the approach of Popović et al. [5] with the uncertainty-only objective and a branch expansion step-size of 5 m empirically set for best performance based on a search over a discrete range of values. We set 60 tree-sampling iterations to obtain the same ~ 76 s average replanning time as required by our replanning routine for a single trajectory. In the random planner, random destinations are sampled in the workspace and a trajectory is generated by connecting them sequentially to the current robot position. We consider 4 waypoints per plan to ensure that the trajectory lengths are fairly comparable to our method.

We conduct 50 simulations in each environment, giving a total of 250 trials. Figure 3 shows how the metrics evolve for each planning method. As expected, in Figure 3b, informed strategies perform better than the random benchmark (blue) as they are guided by planning objectives. In Figure 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 robot pose estimation in the same scenario; confirming that it effectively trades off between gathering information and maintaining good localization given the known landmarks. 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) performs worst as it does not exploit any knowledge of the trajectory dynamics.

Similarly, in Figure 3a, the RIG-tree variant using our Rényi-based (purple) objective provides better localization when compared against planning for map uncertainty reduction only (yellow). This further validates the expected

behaviour of our utility function, as described above, using a different replanning routine. Interestingly, compared to our replanning approach (red), the RIG-tree scores comparatively well on the localization metrics whereas the rate at which the map quality improves is limited. This likely relates to the step-size parameter required by the algorithm, which sets the range of navigation achievable during the mission. By traveling trajectories with shorter steps, the robot re-observes landmarks more frequently at the cost of restricted exploration; leading to a poorer final field reconstruction.

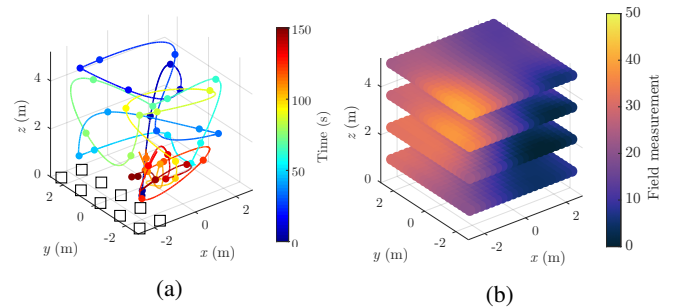


Fig. 4: Example result using our CMA-ES-based replanning framework and our Rényi-based utility function. (a) shows the trajectory traveled by the robot. The squares indicate point landmarks on the ground used for localization. The spheres represent sites where measurements are taken to produce the final field map in (b). By balancing between gathering new information and keeping the landmarks in view, our planner achieves 1.86 lower total field map RMSE compared to the spiral path shown in Figure 1.

An example result using our CMA-ES-based framework and our proposed utility function is shown in Figure 4. Figure 4a confirms that the robot successfully explores the environment while re-visiting known landmarks to stay well-localized. With a total RMSE of 1.11, the map output in this instance (Figure 4b) is 1.86 times more accurate than the one

produced by the naïve spiral path in Figure 1, thus justifying our informative planning strategy.

B. Evaluation of Field Mapping Under Uncertainty

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




Fig. 5: Comparison of field mapping with and without uncertain inputs (UIs) using our CMA-ES-based replanning framework and Rényi-based utility function for a fixed budget of 150 s. The solid lines represent means over 250 trials. The shaded regions show 95% confidence bounds. By accounting for the robot pose uncertainty, our approach achieves more conservative mapping behavior (middle, right) with higher accuracy (left).

Our results are depicted in Figure 5. Note that we omit the map uncertainty metric as the variance scales using the two approaches 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 our modified kernel can handle localization errors to build maps with better consistency and quality for reliable planning. Figure 1 shows a visualization that supports this result.

In summary, Figures 1 and 5 demonstrate that accounting for UIs in the field model improves mapping performance, whereas Figure 3 shows that incorporating the robot pose uncertainty in our utility function yields much better localization. By propagating the pose uncertainty into both the mapping and planning tasks, our proposed active sensing framework offers brings these benefits together to achieve robust mapping behaviour.

C. Proof of Concept

We show our active sensing framework running in real-time on a TurtleBot3 Waffle with an Intel Joule 570x running Ubuntu Linux 16.04 and the Robot Operating System. The experiments are conducted in a known indoor environment using Adaptive Monte Carlo Localization (AMCL)² receiving data from a LDS-01 laser distance scanner. As shown

²wiki.ros.org/amcl

in Figure 6b, for field mapping, a temperature distribution in an empty $2.8\text{ m} \times 2.8\text{ m}$ area within the environment 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 $10\text{ mV}/^\circ\text{C}$.




Fig. 6: Thermal imagery of our experimental setup from side (a) and aerial (b) viewpoints. The robot and radiator are visible. Yellow shades correspond to heated areas mapped using measurements from the on-board temperature sensor.

A 0.4 m resolution grid is set for mapping with UIs. To train the GP, we follow the method in Section VI-B using manually gathered data within the target area to obtain the hyperparameters for the SE kernel. As before, the integral in Equation (4) is estimated using 5 Gauss-Hermite points.

Our uncertainty prediction method is based on localization in a known environment using AMCL (Section V-D). We sub-sample each candidate plan at 2 Hz and estimate the robot pose using a differential drive odometry model and laser scans simulated in the known occupancy map. For the odometry model, variance parameters of 0.2 m^2 are used for Gaussian noise in both rotational and translational motion.

The aim is to show that our framework can map a realistic continuous field using a practical localization system. The initial measurement point is $(0.2\text{ m}, 0.2\text{ m})$ within the corner opposite the radiator. We allocate a planning budget B of 600 s. Following the differential drive model, each plan is piecewise linear as defined by $N = 3$ control waypoints with a constant velocity of 0.26 m/s and temperature measurements sampled at 0.25 Hz . The planning objective is our proposed Rényi-based utility function in Equation (9). Note that field map updates are triggered upon allowing the sensor readings to stabilize between successive measurement points.

Figure 7 summarizes our experiments. As expected, the field map becomes more complete over time and uncertainty decreases as the yellower heated region is discovered in a successful proof of concept implementation. Note that the values of the measurements obtained are higher than that at which the field is initialized (23.64°), i.e., the GP mean. This is due to the effects of heating and diffusion over time, after the training data was collected. Future work will address these issues by investigating map representations that capture temporal, as well as spatial, dynamics.

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




Fig. 7: Experimental results of using our active sensing framework to map the indoor temperature distribution in Figure 6 in a 600 s mission. The three plots on the left depict the trajectories (white lines) and temperature field maps (colored gradients) at different snapshots of the mission at times $t = 100$ s, 350 s, and 600 s. The white circles represent measurement sites, with the large solid one indicating the initial robot position. Yellow shades correspond to hotter regions. The sequence shows that our planner quickly explores the area, successfully detecting the heated corner (bottom-left) where the radiator is located. The curve on the right shows the uncertainty reduction in the field map over time, thus validating our approach. Note that planning time is taken into account.

uses GPs with UIs to propagate the robot pose uncertainty into the model of a target environmental field. For planning, we introduced 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 enables the robot to trade off between exploration and exploitation in a principled way, without relying on any manually tuned parameters.

Our framework was evaluated extensively in simulation. We showed that it achieves more conservative exploratory behavior compared to different planning and mapping strategies, while producing more accurate maps. Experimental validation was performed through a proof of concept deployment, revealing its promise in future applications.

Future work will examine field models with temporal dynamics and efficiency improvements to handle more complex environments. Other interesting research directions involve refining the uncertainty prediction method and its relationship to the Rényi parameter α for more reliable planning.

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*, 2019, under review, arXiv preprint [arXiv:1809.03870](https://arxiv.org/abs/1809.03870).
- [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] S. D. Whitehead and D. H. Ballard, "Active Perception and Reinforcement Learning," *Neural Computation*, vol. 2, no. 4, pp. 409–419, 2008.
- [14] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA: MIT Press, 2006.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] G. Costante, J. Delmerico, M. Werlberger, and P. Valigi, "Exploiting Photometric Information for Planning under Uncertainty," in *Robotics Research*. Springer, 2017, pp. 107–124.
- [19] 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.
- [20] 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.
- [21] N. Hansen, "The CMA evolution strategy: A comparing review," *Studies in Fuzziness and Soft Computing*, vol. 192, no. 2006, pp. 75–102, 2006.
- [22] L. Golshani and E. Pasha, "Rényi entropy rate for Gaussian processes," *Information Sciences*, vol. 180, no. 8, pp. 1486–1491, 2010.
- [23] J. Solà, "Course on SLAM," Barcelona, 2017.
- [24] 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.