



An informative path planning framework for UAV-based terrain monitoring

Marija Popović¹ · Teresa Vidal-Calleja² · Gregory Hitz³ · Jen Jen Chung⁴ · Inkyu Sa⁵ · Roland Siegwart⁴ · Juan Nieto⁴

Received: 25 August 2018 / Accepted: 18 January 2020
© The Author(s) 2020

Abstract

Unmanned aerial vehicles represent a new frontier in a wide range of monitoring and research applications. To fully leverage their potential, a key challenge is planning missions for efficient data acquisition in complex environments. To address this issue, this article introduces a general informative path planning framework for monitoring scenarios using an aerial robot, focusing on problems in which the value of sensor information is unevenly distributed in a target area and unknown a priori. The approach is capable of learning and focusing on regions of interest via adaptation to map either discrete or continuous variables on the terrain using variable-resolution data received from probabilistic sensors. During a mission, the terrain maps built online are used to plan information-rich trajectories in continuous 3-D space by optimizing initial solutions obtained by a coarse grid search. Extensive simulations show that our approach is more efficient than existing methods. We also demonstrate its real-time application on a photorealistic mapping scenario using a publicly available dataset and a proof of concept for an agricultural monitoring task.

Keywords Informative path planning · Aerial robotics · Environmental monitoring · Remote sensing

M. Popović, G. Hitz, I. Sa: Autonomous Systems Lab., ETH Zürich, Zürich, Switzerland at the time of this work.

✉ Marija Popović
m.popovic@imperial.ac.uk

Teresa Vidal-Calleja
teresa.vidalcallega@uts.edu.au

Gregory Hitz
gregory.hitz@sevensesense.ch

Jen Jen Chung
jenjen.chung@mavt.ethz.ch

Inkyu Sa
inkyu.sa@csiro.au

Roland Siegwart
rsiegwart@ethz.ch

Juan Nieto
nietoj@ethz.ch

¹ Smart Robotics Lab., Imperial College London, London, UK

² Centre for Autonomous Systems, University of Technology, Sydney, Australia

³ Sevensesense, Zürich, Switzerland

⁴ Autonomous Systems Lab., ETH Zürich, Zürich, Switzerland

1 Introduction

Autonomous mobile robots are increasingly employed to gather valuable scientific data about the Earth. In the past several decades, rapid technological advances have unlocked their potential as a flexible, cost-efficient tool enabling monitoring at unprecedented levels of resolution and autonomy. In many emerging aerial (Ezequiel et al. 2014; Vivaldini et al. 2016; Colomina and Molina 2014) and aquatic (Hitz et al. 2014, 2017; Girdhar and Dudek 2015) applications, these devices are replacing traditional data acquisition campaigns based on static sensors, manual sampling, or conventional manned platforms, which can be unreliable, costly, and even dangerous (Dunbabin and Marques 2012; Manfreda et al. 2018).

The era of robotics-based monitoring has opened many interesting areas of research. A key challenge arises in that practical devices are subject to a finite quantity of sensing resources, such as energy, time, or travel distance, which limits the number of measurements that can be collected. Therefore, paths need to be planned to maximize the

⁵ Dynamic Platforms, Robotics and Autonomous Systems Group, Data61, CSIRO, Brisbane, Australia

information gathered about an unknown environment while satisfying the given budget constraint. This is known as the *informative path planning (IPP)* problem, which is the subject of much recent work. Despite recent advances in this field, most real robot systems today still perform data acquisition in a passive manner, e.g., using coverage-based planning (Galceran and Carreras 2013), as current IPP solutions tend to be limited to specific platforms or application domains.

To address this, our work introduces a general framework for surveying terrain characteristics using an aerial robot. We present an IPP strategy suitable for *adaptive* scenarios, which require modifying plans online based on new measurements to learn and focus on targeted objects or regions of interest, e.g., to find cars in a parking lot or map high-temperature areas. This setup involves several closely related research challenges. First, the dense visual information received from different altitudes needs to be integrated into a compact probabilistic map. Second, using the map, the planning unit must search for informative trajectories in the large 3-D space above the monitored area, which poses a complex optimization problem. During this procedure, a crucial aspect is trading off between sensor resolution and Field of View (FoV), while accounting for the platform-specific constraints and adaptivity requirements. Finally, the integrated system is required to be efficient for online mapping and planning in a wide variety of scenarios. In this paper, we cater for these needs by developing a modular IPP framework for aerial robots in general environmental monitoring applications.

1.1 Contributions

Overall, this article corresponds to a major extension of the methods introduced in the authors' preliminary conference works. The journal version brings together our previous contributions in online informative planning (Popović et al. 2017a) and multiresolution mapping (Popović et al. 2017b) with additional explanations and simulations to consolidate the integrated system along with open-source software.

Building upon our developments in planning and mapping, we present an IPP framework capable of mapping either discrete or continuous target variables on a terrain using a probabilistic altitude-dependent sensor. A key feature of the approach is that it is modular, with interfaces and computational requirements that are easily adapted for a given monitoring scenario. During a mission, the planner uses the terrain maps built online to generate trajectories in continuous 3-D space for maximized information gain. This is achieved in a computationally efficient manner by exploiting recursive Bayesian mapping methods and an optimization strategy for IPP with an informed initialization procedure. Our method was tested extensively in simulation and its online and integration capabilities were validated by mapping a publicly available dataset using an Unmanned

Aerial Vehicle (UAV) equipped with an image-based classifier. Additionally, a proof of concept field deployment is presented in an agricultural monitoring application with real-time requirements.

The core contributions of this work are:

1. An informative planning framework that is applicable for mapping either discrete or continuous variables on a given terrain using probabilistic sensors for data acquisition in online settings with adaptivity requirements by learning targeted objects or regions of interest.
2. The extensive evaluation of our framework in simulation, along with results from a publicly available dataset and field experiments to demonstrate its performance.
3. The release of our implementation as an open-source software package¹ for usage and further development by the community.

The remainder of this paper is organized as follows. Section 2 discusses prior studies relevant to our work. We define the IPP problem in Sect. 3 and detail our proposed methods for mapping and planning in Sects. 4 and 5, respectively. Our experimental results are reported in Sects. 6 and 7. Finally, in Sect. 8, we conclude with an outlook towards future work.

2 Related work

A large and growing body of literature exists for autonomous information gathering problems. Recently, this field has attracted considerable interest in the context of robotics-based environmental monitoring (Dunbabin and Marques 2012) for a wide variety of applications, including aerial surveillance (Colomina and Molina 2014; Vivaldini et al. 2016), aquatic monitoring (Hitz et al. 2014, 2017), and infrastructure inspection (Ezequiel et al. 2014; Bircher et al. 2016). This section overviews recent work based on two main research streams: (1) methods for environmental modeling and (2) algorithms for informative planning, or efficient data acquisition.

2.1 Environment mapping

In data gathering scenarios, a model of the environment is fundamental to capture the target variable of interest. Occupancy grids are the most commonly-used representation for spatial sensing with uncorrelated measurements (Elfes 1989). This type of model is suitable for active classification problems with discrete labels, such as occupancy mapping (Charrow et al. 2015) and semantic segmentation (Berrio

¹ Available at: <http://github.com/ethz-asl/tmplanner>.

et al. 2017), and offers relatively high computational efficiency.

However, many natural phenomena exhibit complex interdependencies where the assumption of independent measurements does not hold. A popular Bayesian technique for handling such relationships is using Gaussian Processes (GPs) (Rasmussen and Williams 2006). For IPP, they have been applied in various scenarios (Hollinger and Sukhatme 2014; Hitz et al. 2017; Binney and Sukhatme 2012; Wei et al. 2016) to collect data accounting for map structure and uncertainty. This framework permits using different kernel functions to express data relations within the environment (Singh et al. 2010) and approximations to scale to large datasets (Rasmussen and Williams 2006).

For terrain monitoring, the main issue with applying GPs directly is the escalating computational load as dense imagery data accumulate over time. Moreover, limited studies have addressed data fusion using images at varying resolutions and noise levels. In our prior work, we addressed this by introducing a multiresolution GP-based mapping approach for IPP based on a Bayesian filtering technique, inspired by Vidal-Calleja et al. (2014). This method replaces the complexity of conventional GP regression with constant processing time in the number of measurements and cubic scaling only with the size of the fixed terrain map. Our current work extends these ideas by presenting a general framework that can handle both non-correlated and correlated monitored variables and by demonstrating its integration with a practical altitude-dependent sensor.

2.2 Informative planning

We also distinguish between (i) non-adaptive and (ii) adaptive planning strategies. Non-adaptive approaches, e.g., coverage methods (Galceran and Carreras 2013), explore an environment using a sequence of pre-determined actions. Adaptive approaches (Hitz et al. 2017; Girdhar and Dudek 2015; Lim et al. 2015; Hollinger et al. 2013) allow plans to change as information is collected, making them suitable for planning based on targeted application-specific interests.

In its most general form, the data gathering task amounts to one of sequential decision-making under uncertainty, which can be expressed as a Partially Observable Markov Decision Process (POMDP) (Kaelbling et al. 1998). Unfortunately, despite substantial progress in recent years (Chen et al. 2016; Kurniawati et al. 2008), solving large-scale POMDP models remains an open challenge. For data gathering, a major issue with practical algorithms is that they do not generalize well for large state spaces or over long horizons (Lim 2015). Although approximate Markov Decision Process (MDP)-based solvers have been successfully applied for online planning, their performance is limited in terrain monitoring setups since they cannot exploit the ability of

UAVs to gather multiresolution information from different altitudes, e.g., to adaptively find objects of interest (Arora et al. 2018). The computational and representational issues associated with such methods motivates more efficient IPP-based solutions.

The NP-hard sensor placement problem addresses selecting the most informative measurement sites in a static setting subject to cardinality constraints (Krause et al. 2008). The related IPP problem builds upon this task by maximizing the information gathered by a mobile platform traversing a path while subject to finite resources. A number of general algorithms for IPP operate by performing combinatorial optimization over a discrete grid. The seminal work of Chekuri and Pál (2005) presents a recursive greedy algorithm providing a sub-logarithmic approximation in sub-exponential time. Singh et al. (2009) apply these ideas to environmental sensing problems and extend them to multi-robot settings. More recently, branch and bound techniques have been proposed to improve computational efficiency using monotonic objective functions (Binney and Sukhatme 2012). In 2-D scenarios, an alternative strategy is to represent the robot workspace as decomposed cells (Cai and Ferrari 2009) or a connectivity tree (Ferrari and Cai 2009). These methods (Cai and Ferrari 2009; Binney and Sukhatme 2012) have been shown to outperform coverage-based planning strategies (Choset 2001; Galceran and Carreras 2013) and are applicable in cluttered environments.

However, discrete solvers are typically limited in resolution and scale exponentially with the problem instance. To reduce the computational overhead, we follow previous approaches (Bircher et al. 2016; Hitz et al. 2014) in using a method with a limited look-ahead to generate plans within a fixed horizon of the total budget. This approach allows us to perform incremental replanning, as required for adaptive settings, by trading off guarantees on exploration optimality outside the planning horizon.

Continuous-space planning strategies offer better scalability in comparison to discrete methods by leveraging Rapidly-exploring Random Tree (RRT)-like sampling-based methods (Hollinger and Sukhatme 2014) or splines (Vivaldini et al. 2016; Hitz et al. 2017; Charrow et al. 2015; Morere et al. 2017) directly in the robot workspace. As in our prior work (Popović et al. 2017a), we follow the latter approaches in defining smooth polynomial trajectories (Richter et al. 2013) which are optimized globally for an information objective. Our spline optimization problem setup most closely resembles the ones studied by Hitz et al. (2017) and Morere et al. (2017); however, our strategy differs in that it uses an informed initialization procedure to obtain faster convergence.

Specifically, we follow a two-step approach, which exploits a greedy, grid-based solution to the resource-constrained IPP problem to initialize an evolutionary optimization rou-

tine. Our grid search procedure closely resembles frontier-based strategies commonly used for exploration in cluttered environments (Yamauchi 1998). However, a key difference is that our method does not rely on the internal geometry of the environment as the camera FoV is unobstructed by obstacles. For the optimization step, an alternative approach to our global evolutionary strategy is to generate control actions for the robot using gradient ascent on an information objective (Schwager et al. 2017; Julian et al. 2011; Grocholsky et al. 2003). Such techniques would be useful when the initial grid-based solution is close to a trajectory with high information quality, but may be susceptible to local minima. This issue can be addressed by applying algorithms based on random walks, local roadmaps (Lu et al. 2014; Bellini et al. 2014), or two-stage optimization strategies (Charrow et al. 2015; Rastgoftar et al. 2018).

The motivation behind our approach is shared by the studies of Jadidi et al. (2016, 2018) in robotic exploration and information gathering. In contrast, however, we address the problem of terrain monitoring, considering a 3-D robot workspace, instead of 2-D occupancy mapping and environmental mapping with a point-based sensor.

Within this context, several recent works have examined applications similar to ours. Notably, Sadat et al. (2015) devise an adaptive planner for aerial coverage problems in which regions of interest are non-uniformly distributed. Their method, however, assumes discrete viewpoints and does not support probabilistic data acquisition. To address this, Bellini et al. (2014) present an information gradient-based control law for active target classification in cluttered environments. More similarly to our approach, Arora and Scherer (2017) introduce an efficient near-optimal algorithm that provides anytime solutions in adaptive scenarios. As further discussed by Arora et al. (2018), their setup considers using a multiresolution sensor to gather targeted information about specific objects as they are detected. A key difference is that our method also caters for probabilistic variations in sensor *noise* for data fusion at different altitudes, and can accommodate adaptively mapping continuous variables, e.g., temperature, as well as discrete objects, e.g., cars.

3 Problem statement

Our setup focuses on efficient data-gathering strategies for an aerial robot operating above a terrain. The aim is to maximize the information collected about the environment, while respecting resource constraints, such as energy, time, or distance budgets. Formally, this is known as the IPP problem, which is defined as follows. We seek an optimal trajectory ψ^* in the space of all continuous trajectories Ψ for maximum gain in some information-theoretic measure:

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

The function $\text{MEASURE}(\cdot)$ obtains a finite set of measurements along trajectory ψ in the 3-D space above the environment, and $\text{COST}(\cdot)$ provides the corresponding cost, which cannot exceed a predefined budget B . The operator $I(\cdot)$ defines the information objective quantifying the utility of the acquired measurements.

4 Mapping approach

In this section, we propose a new mapping framework for terrain monitoring applications. The generic structure of our system setup is depicted in Fig. 1. As shown, our framework is capable of mapping either discrete or continuous variables based on measurements extracted from a sensing unit, e.g., a depth or multispectral camera. For a particular problem setup, the map representation can be selected depending on the type of data received. During a mission, the planner uses the terrain maps built online to optimize continuous trajectories for maximum gain in an information metric reflecting the mission aim. A key aspect of our architecture is its generic formulation, which enables it to adapt to any surface mapping scenario, e.g., elevation (Colomina and Molina 2014), pipe thickness (Vidal-Calleja et al. 2014), gas concentration (Marchant and Ramos 2014), spatial occupancy (O’Callaghan and Ramos 2012), seismic hazards (Gao et al. 2017), post-disaster assessment (Ezequiel et al. 2014), signal strength (Hollinger and Sukhatme 2014), etc.

In Sects. 4.1 and 4.2, we present methods for map representation for monitoring discrete and continuous targets, respectively, as the basis of our framework. In Sect. 5, these concepts are used to formulate the objective function, and we describe our adaptive planning scheme.

4.1 Discrete variable mapping

We study the task of monitoring a discrete variable as an active classification problem. The terrain environment \mathcal{E} is discretized and represented using a 2-D occupancy map \mathcal{X} (Elfes 1989), where each grid cell is associated with an independent Bernoulli random variable indicating the probability of target occupancy (e.g., presence of weed on a farmland). Measurements are taken with a downwards-looking sensor providing inputs to a data processing unit, from which discrete classification labels are obtained. At time t , for each observed cell $\mathbf{x}_i \in \mathcal{X}$ within the FoV from a UAV pose \mathbf{p} above the terrain, a log likelihood update is performed given an observation z :




Fig. 1 System diagram showing the key elements of our IPP framework. A map of the target environment is built using measurements extracted from a sensor data stream. The dashed lines indicate that the map representation can be either discrete or continuous. At a particular

time instant, the map knowledge is used by the planning unit to find the most useful trajectories for data collection, starting at the current pose. These are then executed by the robot, allowing for subsequent map updates in a closed-loop manner




Fig. 2 Sensor model for a typical camera-based binary classifier operating above a terrain. The blue and orange curves depict the probability of observing label ‘1’ for a map cell containing ‘1’ or ‘0’, respectively, i.e., $p(z | \mathbf{x}_i, \mathbf{p})$. As altitude increases, the curves approach unknown classification probability (0.5)

$$L(\mathbf{x}_i | z_{1:t}, \mathbf{p}_{1:t}) = L(\mathbf{x}_i | z_{1:t-1}, \mathbf{p}_{1:t-1}) + L(\mathbf{x}_i | z_t, \mathbf{p}_t), \tag{2}$$

where $L(\mathbf{x}_i | z_t, \mathbf{p}_t)$ denotes the altitude-dependent inverse sensor model log-likelihood capturing the classification output.

As an example, Fig. 2 shows the sensor model for a hypothetical camera-based binary classifier labeling observed cells as ‘1’ (occupied by target) or ‘0’ (target-free). For each class, curves are defined to account for poorer classification with lower-resolution measurements taken at higher altitudes. In practice, these curves can be obtained through a Monte Carlo-type accuracy analysis of raw classifier data

by averaging the number of true and false positives (blue and orange curves, respectively) recorded at different altitudes.

The described approach can be easily extended to mapping multiple class labels by maintaining layers of occupancy maps for each, as demonstrated in Sect. 6.2.

4.2 Continuous variable mapping

To monitor a continuous variable, our framework leverages a more sophisticated mapping method using GPs to encode spatial correlations common in environmental distributions. We use a GP to initialize a recursive filtering procedure with probabilistic sensors at potentially different resolutions. This approach replaces the computational burden of applying GPs directly with constant processing time in the number of measurements. We describe our method for creating prior maps before detailing the Bayesian data fusion technique.

4.2.1 Gaussian processes

A GP is used to model spatial correlations on the terrain in a probabilistic and non-parametric manner (Rasmussen and Williams 2006). The target variable for mapping is assumed to be a continuous function in 2-D space: $\zeta : \mathcal{E} \rightarrow \mathbb{R}$. Using the GP, a Gaussian correlated prior is placed over the function space, which is fully characterized by the mean $\boldsymbol{\mu} = \mathbb{E}[\zeta]$ and covariance $\mathbf{P} = \mathbb{E}[(\zeta - \boldsymbol{\mu})(\zeta^\top - \boldsymbol{\mu}^\top)]$ as $\zeta \sim \mathcal{GP}(\boldsymbol{\mu}, \mathbf{P})$, where $\mathbb{E}[\cdot]$ denotes the expectation operator.

Given a pre-trained kernel $K(\mathcal{X}, \mathcal{X})$ for a fixed-size terrain discretized at a certain resolution with a set of n locations $\mathcal{X} \subset \mathcal{E}$, we first specify a finite set of new prediction points $\mathcal{X}^* \subset \mathcal{E}$ at which the prior map is to be inferred. For unknown environments, as in our setup, the values at $\mathbf{x}_i \in \mathcal{X}$ are ini-

tialized uniformly with a constant prior mean. For known environments, the GP can be trained from available data and inferred at the same or different resolutions. The covariance is calculated using the classic GP regression equation (Reece and Roberts 2013):

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

where \mathbf{P} is the posterior covariance, \mathbf{I}_n is the $n \times n$ identity matrix, σ_n^2 is a hyperparameter representing signal noise variance, and $K(\mathcal{X}^*, \mathcal{X})$ denotes cross-correlation terms between the predicted and initial locations.

The kernel $K(\cdot)$ determines the generalization properties of the GP model, and is chosen to describe the characteristics of ζ . To describe environmental phenomena, we suggest choosing from among a number of well-known kernel functions common in geostatistical analysis, e.g., the squared exponential or Matérn functions. The free parameters of the kernel function, called hyperparameters, control relations within the GP. These values can be optimized using various methods (Rasmussen and Williams 2006) to match the properties of ζ by training on multiple maps obtained previously at the required resolution.

Once the correlated prior map $p(\zeta | \mathcal{X})$ is determined, independent noisy measurements at variable resolutions are fused as described in the following section.

4.2.2 Sequential data fusion

A key component of our framework is a map update procedure based on recursive filtering. Given a uniform mean and the spatial correlations captured by Eq. (3), the map $p(\zeta | \mathcal{X}) \sim \mathcal{GP}(\boldsymbol{\mu}^-, \mathbf{P}^-)$ is used as a prior for fusing new sensor measurements.

Let $\mathbf{z} = [z_1, \dots, z_m]^\top$ denote new m independent measurements received at points $[\mathbf{x}_1, \dots, \mathbf{x}_m]^\top \subset \mathcal{X}$ modeled assuming a Gaussian sensor as $p(z_i | \zeta_i, \mathbf{x}_i) = \mathcal{N}(\mu_{s,i}, \sigma_{s,i})$. To fuse the measurements \mathbf{z} with the prior $p(\zeta | \mathcal{X})$, we use the maximum *a posteriori* estimator:

$$\underset{\zeta}{\operatorname{argmax}} p(\zeta | \mathbf{z}, \mathcal{X}). \quad (4)$$

The Kalman Filter (KF) update equations are applied directly to compute the posterior density $p(\zeta | \mathbf{z}, \mathcal{X}) \propto p(\mathbf{z} | \zeta, \mathcal{X}) \times p(\zeta | \mathcal{X}) \sim \mathcal{GP}(\boldsymbol{\mu}^+, \mathbf{P}^+)$ (Reece and Roberts 2013):

$$\boldsymbol{\mu}^+ = \boldsymbol{\mu}^- + \mathbf{K}\mathbf{v}, \quad (5)$$

$$\mathbf{P}^+ = \mathbf{P}^- - \mathbf{K}\mathbf{H}\mathbf{P}^-, \quad (6)$$

where $\mathbf{K} = \mathbf{P}^- \mathbf{H}^\top \mathbf{S}^{-1}$ is the Kalman gain, and $\mathbf{v} = \mathbf{z} - \mathbf{H}\boldsymbol{\mu}^-$ and $\mathbf{S} = \mathbf{H}\mathbf{P}^- \mathbf{H}^\top + \mathbf{R}$ are the measurement and covariance

innovations. \mathbf{R} is a diagonal $m \times m$ matrix of altitude-dependent variances $\sigma_{s,i}^2$, associated with each measurement z_i , and \mathbf{H} is an $m \times n$ matrix denoting a linear measurement model that intrinsically selects part of the state $\{\zeta_1, \dots, \zeta_m\}$ observed through \mathbf{z} . The information to account for variable-resolution measurements is incorporated according to the measurement model \mathbf{H} in a simple manner as detailed in the following section.

The constant-time updates in Eqs. (5) and (6) are repeated every time new data are registered. Note that, as all models are linear in this case, the KF update produces the optimal solution. Moreover, this approach is agnostic to the type of sensor used as it permits fusing heterogeneous data into a single grid map.

4.2.3 Altitude-dependent sensor model

As an example, we detail an altitude-dependent sensor model for a downward-facing camera used to take measurements of a terrain, e.g., a farmland or disaster site. In contrast with the pure classification case in Sect. 4.1, our model needs to express uncertainty with respect to a continuous target distribution. To do this, we consider that the visual data degrades with altitude in two ways: (i) noise and (ii) resolution. The proposed model accounts for these issues in a probabilistic manner as follows.

We assume an altitude-dependent Gaussian sensor noise model. For each observed point $\mathbf{x}_i \in \mathcal{X}$, the camera provides a measurement z_i capturing the target field ζ_i as $\mathcal{N}(\mu_{s,i}, \sigma_{s,i})$, where $\sigma_{s,i}^2$ is the noise variance expressing uncertainty in z_i . To account for lower-quality images taken with larger camera footprints, $\sigma_{s,i}^2$ is modeled as increasing with the UAV altitude h using:

$$\sigma_{s,i}^2 = a(1 - e^{-bh}), \quad (7)$$

where a and b are positive constants.

Figure 3 illustrates the sensor noise model used to evaluate our setup in Sect. 6 which represents a camera. The measurements z_i denote the levels of the continuous variable being surveyed, e.g., green biomass level or temperature. This model is designed so that the quality of the sensor data decreases at higher altitudes, according to requirements of the terrain monitoring problem. As for the discrete classifier in Sect. 4.1, in practice, the parameters of such a model can be obtained by analyzing how the sensor behaves at different altitudes using previously acquired datasets.

We define altitude envelopes corresponding to different resolution scales with respect to the initial points \mathcal{X} on the map. This is motivated by the fact that the Ground Sample Distance (GSD) ratio (in m/px) depends on the altitude of the sensor and its fixed intrinsic resolution. To handle data received from variable altitudes, adjacent \mathbf{x}_i are indexed by




Fig. 3 Inverse sensor noise model for a camera providing measurements as $\mathcal{N}(\mu_{s,i}, \sigma_{s,i}^2)$ with $a = 0.2$, $b = 0.05$ in Eq. (7). The uncertainty $\sigma_{s,i}^2$ increases with h to represent degrading image quality. The dotted line at $h = 10$ m indicates the altitude above which image resolution scales down by a factor of 2

a single sensor measurement z_i through the measurement model \mathbf{H} . At lower altitudes (higher GSDs, corresponding to the maximum mapping resolution in \mathcal{X}), \mathbf{H} is simply used to select the part of the state observed with a scale of 1. However, at higher altitudes (lower GSDs), the elements of \mathbf{H} used to map multiple z_i to a single z_i are scaled by the square inverse of the resolution scaling factor s_f . Note that the fusion procedure described in Sect. 4.2.2 is always performed at the maximum mapping resolution as defined by \mathcal{X} , so that the proposed model \mathbf{H} considers low-resolution measurements as a scaled average of the high-resolution map.

Figure 4 consolidates our strategy with an example. The map in (d) depicts the ground truth corresponding to a field distribution on a terrain. Mapping is performed at the resolution shown in (a), where the grid cells correspond to the locations in \mathcal{X} . The maps in (e)–(h) demonstrate sequentially fusing two measurements taken at different altitudes (middle row) into a single probabilistic map (bottom row), i.e., (g) and (h) visualize the results of fusing first (e) then (f), respectively, assuming map initialization with a uniform mean. The plots in (b) and (c) (top row) schematize the UAV poses from which the measurements were registered, with the red grids indicating their resolutions. In (c), the scaling factor is $s_f = 0.5$, such that 4 locations in \mathcal{X} (black cells) map to a single measurement (red cell). By inspecting the final field map in (h), upon fusing (f), it can be seen the off-center values are more widely diffused compared to those in the center, where the higher-quality measurement in (e) was taken, as expected.

5 Planning approach

This section details our planning scheme for terrain monitoring. As depicted in Fig. 1, we generate fixed-horizon plans to




Fig. 4 Overview of our strategy for continuous variable mapping. A ground truth field is shown in (d). e, f depict measurements taken from 9 m and 20 m altitudes, with different resolutions visible in the projected camera FoVs (pyramids) in (b) and (c). The resolution scales (red) are defined by factors s_f of 1 and 0.5 with respect to the mapping resolution in (a), respectively. g, h illustrate maps resulting from fusing the data sequentially. The variable diffusion effects show that our method can accommodate noisy multiresolution measurements and capture spatial correlations (Color figure online)

maximize an objective. To do this efficiently, an evolutionary technique is applied to optimize trajectories initialized by a 3-D grid search in the UAV workspace. We first describe our approach to parameterizing trajectories, then detail the algorithm itself.

5.1 Trajectories

A polynomial trajectory ψ is represented by a sequence of N 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 (Richter et al. 2013). The first waypoint \mathbf{c}_1 is clamped as the initial UAV position. As discussed in Sect. 3, the function $\text{MEASURE}(\cdot)$ in Eq. (1) is defined by computing the spacing of measurement sites along ψ given a constant sensor frequency and the traveling speed of the UAV along the trajectory.

5.2 Algorithm

A fixed-horizon approach is used to plan adaptively as data is collected. During a mission, we alternate between replanning and execution until the elapsed time t exceeds a budget B .

Our replanning approach consists of two stages and is shown in Algorithm 1. First, an initial trajectory, defined by N fixed control waypoints \mathcal{C} , is derived through a coarse grid search (Lines 3–7) in the 3-D workspace. This step proceeds sequentially, selecting points in a greedy manner based on a chosen utility function $I(\cdot)$, so that a rough solution can quickly be obtained. Then, the trajectory is refined using Eq. (1) to maximize the information objective. In this step, we employ a generic evolutionary routine (Line 8) to optimize the set of control waypoints \mathcal{C} .

In Algorithm 1, \mathcal{Z} symbolizes a general model of the environment \mathcal{E} capturing either a discrete or continuous target variable of interest. The following sections outline the key steps of the algorithm before discussing possible ways of defining the objective for informative planning, including in scenarios with adaptivity requirements that require learning and focusing on objects or regions of interest online as they are discovered.

Algorithm 1 REPLAN_PATH procedure

Input: Current environment model \mathcal{Z} , initial position \mathbf{c}_1 , number of control waypoints N , lattice points \mathcal{L}

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

```

1:  $\mathcal{Z}' \leftarrow \mathcal{Z}$  // Create a local copy of the map.
2:  $\mathcal{C} \leftarrow \mathbf{c}_1$ ;  $\mathbf{c}_{\text{prev}} \leftarrow \mathbf{c}_1$  // Initialize control points.
3: while  $|\mathcal{C}| \leq N$  do
4:    $\mathbf{c}^* \leftarrow \underset{\mathbf{c} \in \mathcal{L}}{\text{argmax}} \frac{I(\mathbf{c})}{\text{COST}(\mathbf{c}, \mathbf{c}_{\text{prev}})}$  // Find next-best point.
5:    $\mathcal{Z}' \leftarrow \text{PREDICT\_MEASUREMENT}(\mathcal{Z}', \mathbf{c}^*)$  // From this point.
6:    $\mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}^*$  // Add to trajectory solution.
7:    $\mathbf{c}_{\text{prev}} \leftarrow \mathbf{c}^*$ 
8:  $\mathcal{C} \leftarrow \text{CMAES}(\mathcal{C}, \mathcal{Z})$  // Optimize control points using Eq. (1).
```

5.3 3-D grid search

The first step of the replanning procedure (Lines 3–7 of Algorithm 1) supplies an initial solution for the optimization step in Sect. 5.4. To achieve this, the planner performs a 3-D grid search based on a coarse multiresolution lattice \mathcal{L} in the robot workspace (Fig. 5). A low-accuracy solution neglecting sensor dynamics is obtained efficiently by using the points in \mathcal{L} to represent candidate measurement sites and assuming constant velocity travel. Unlike in frontier-based exploration commonly used in cluttered environments (Charrow et al. 2015), selecting goal measurement sites based on map boundaries is not applicable in our setup. Instead, we conduct a sequential greedy search for N waypoints (Line 3), where the next-best point \mathbf{c}^* (Line 4) is found by evaluating the incremental information gain rate given the chosen utility definition $I(\cdot)$ over \mathcal{L} . This term represents the information objective and is discussed in detail in Sect. 5.5. Since the number of waypoints N is specified as a constant, note that our approach in this step closely resembles the sequen-

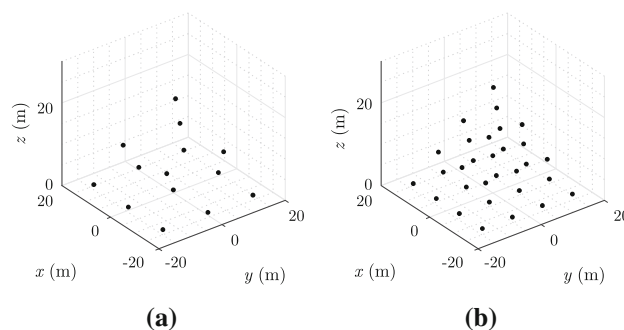


Fig. 5 Visualizations of **a** 14-point and **b** 30-point 3-D lattice grids \mathcal{L} for obtaining an initial trajectory solution in a $40 \text{ m} \times 40 \text{ m} \times 30 \text{ m}$ space. The point density can be chosen to trade off between solution accuracy and computational efficiency in the grid search. Note that the points are sparser at the top due to increasing FoV

tial greedy algorithm (Krause et al. 2008) and generalized cost-benefit greedy algorithm (Zhang and Vorobeychik 2016) for submodular function optimization with cardinality constraints, with the requirement that travel time of the output trajectory defined by the points must lie within the budget B . However, the optimization subproblem in Algorithm 1 fundamentally differs from cardinality constrained submodular maximization in that our objective function is *not* submodular.

Importantly, the prediction step assumes that no prior knowledge about the value of future measurements is available. The search is thus conditioned on the *most likely estimate* of the current field map, i.e., considering that the maximum probability distribution of the map state will be re-observed. For discrete mapping, this involves partitioning the occupancy grid cells with $p(\mathbf{x}_i) \geq 0.5$ as being occupied and those with $p(\mathbf{x}_i) < 0.5$ as being free in order to identify which of the two classification curves in Fig. 2 to apply for the map update procedure (Eq. (2)). For continuous variable mapping, the most likely estimate of the map simply corresponds to the current mean μ of the field distribution and the update is performed using the Bayesian data fusion technique (Eqs. (5) and (6)).

For each \mathbf{c}^* , a fused measurement is simulated in \mathcal{Z} using the relevant update strategy (Line 5). This point is then added to the initial trajectory solution (Line 6).

As depicted in Fig. 5, the length scales of \mathcal{L} can be defined based on the computational resources available and the level of accuracy desired; the denser grid in Fig. 5b procures better initial solutions at the expense of longer evaluation times.

5.4 Optimization

The second step (Line 8 of Algorithm 1) optimizes the coarse grid search solution for the control waypoints \mathcal{C} using Eq. (1). This is done by computing $I(\cdot)$ for a sequence of measurements taken along the trajectory, as defined in Sect. 5.5.

Thereby, an informed initialization procedure is used to speed up the convergence of the optimizer. Note that this step is agnostic to the optimization method considered; in our specific approach, we apply the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), discussed below. In Sect. 6.1.2, we evaluate our choice by comparing the use of different routines.

The CMA-ES is a generic global optimization routine based on the concepts of evolutionary algorithms which has been successfully applied to high-dimensional, nonlinear, non-convex problems in the continuous domain. As an evolutionary strategy, the CMA-ES operates by iteratively sampling candidate solutions according to a multivariate Gaussian distribution in the search space. Further details, including a convergence analysis, are provided in the in-depth review by Hansen (2006).

Our choice of optimization method is motivated by the nonlinearity of the objective space in Eq. (1) as well as by previous results (Popović et al. 2017a, b; Hitz et al. 2017). We initialize the mean solution with the grid search result, constraining \mathbf{c}_1 to coincide with the current robot pose, as described above. In addition, a coordinate-wise boundary handling algorithm (Hansen 2009) is applied to constrain the measurement points to lie within a feasible cubical volume of the workspace above the terrain. For optimization, we use the strategy internal parameters proposed by Hansen (2006), and tune only the population size (number of offspring in the search), co-ordinate wise initial standard deviation (step size, representing the distribution spread), and maximum number of iterations based on application-specific requirements.

5.5 Utility definition

The utility, or information gain, function $I(\cdot)$ is critical as it encapsulates the specific interests for data-driven planning (Eq. (1)). This section discusses possible ways of quantifying the value of new sensor measurements with respect to the proposed map representations. We first introduce utility functions for a pure exploration scenario, in which information gain depends only on the map uncertainty. Sect. 5.5.1 then discusses objectives for missions with adaptivity requirements, where the aim is to discover and focus on targeted regions of interest in the environment.

We examine definitions of $I(\cdot)$ for evaluating the exploratory value of a measurement from a pose \mathbf{p} (Line 4 of Algorithm 1). Note that, above, \mathbf{c} denotes a control waypoint parameterizing a polynomial trajectory, whereas here \mathbf{p} is a generic pose from where the measurement is registered. In particular, in a pure non-targeted exploration scenario, we consider maximizing the reduction of Shannon's entropy H in the map \mathcal{X} :

$$I(\mathbf{p}) = H(\mathcal{X}^-) - H(\mathcal{X}^+), \quad (8)$$

where the superscripts denote the prior and posterior maps given a measurement taken from \mathbf{p} .

In the discrete variable scenario, the value of H for the occupancy map \mathcal{X} is obtained by simply summing over the entropy values of all cells $\mathbf{x}_i \in \mathcal{X}$, assuming their independence:

$$H(\mathcal{X}) = - \sum_{\mathbf{x}_i \in \mathcal{X}} p(\mathbf{x}_i) \log p(\mathbf{x}_i) + (1 - p(\mathbf{x}_i)) \log (1 - p(\mathbf{x}_i)), \quad (9)$$

where $p(\mathbf{x}_i)$ indicates the probability of occupancy at \mathbf{x}_i .

In the continuous variable scenario, however, calculating H involves the determinant of the covariance matrix \mathbf{P} of the GP model (Rasmussen and Williams 2006). We avoid this computationally expensive step by instead maximizing the decrease in the matrix trace, which is equivalent to minimizing the criterion for A-optimal design, and only measures the total variance of the map cells (Sim and Roy 2005):

$$I(\mathbf{p}) = \text{Tr}(\mathbf{P}^-) - \text{Tr}(\mathbf{P}^+), \quad (10)$$

where $\text{Tr}(\cdot)$ denotes the trace of a matrix.

Note that Eq. (8) defines $I(\cdot)$ for a single measurement site \mathbf{p} . To determine the utility of a complete trajectory ψ , the same principles can be applied by fusing a sequence of measurements iteratively and computing the overall information gain.

5.5.1 Adaptivity for regions of interest

We also study an adaptive planning setup where the aim is to map targeted objects or areas of interest, such that the objective depends on the values of the measurements taken in addition to their location. This property is very valuable for practical monitoring applications, such as finding function extrema (Marchant and Ramos 2014), classifying level sets (Hitz et al. 2014), or focusing on specific value ranges. To this end, Eq. (8) is modified so that the elements mapping to the value of each cell $\mathbf{x}_i \in \mathcal{X}$ are excluded from the objective computation, provided they do not satisfy the requirement which defines interest-based planning.

As an example, this work considers a mission where the aim is to focus specifically on regions of interest which have higher values of the latent target parameter, e.g., areas of high vegetation cover in a field. A threshold is applied to separate the (a) "interesting" (above) and (b) "uninteresting" (below) parameter value range into complementary subsets (a) \mathcal{X}_I and (b) \mathcal{X}_- of all points \mathcal{X} in the map. The partitioning strategy is described below. The main idea is to selectively include only the points \mathcal{X}_I in calculating the information value of potential measurements in Eq. (8), i.e., the utility associated with lower values in the points $\mathcal{X}_- = \mathcal{X} \setminus \mathcal{X}_I$ is ignored.

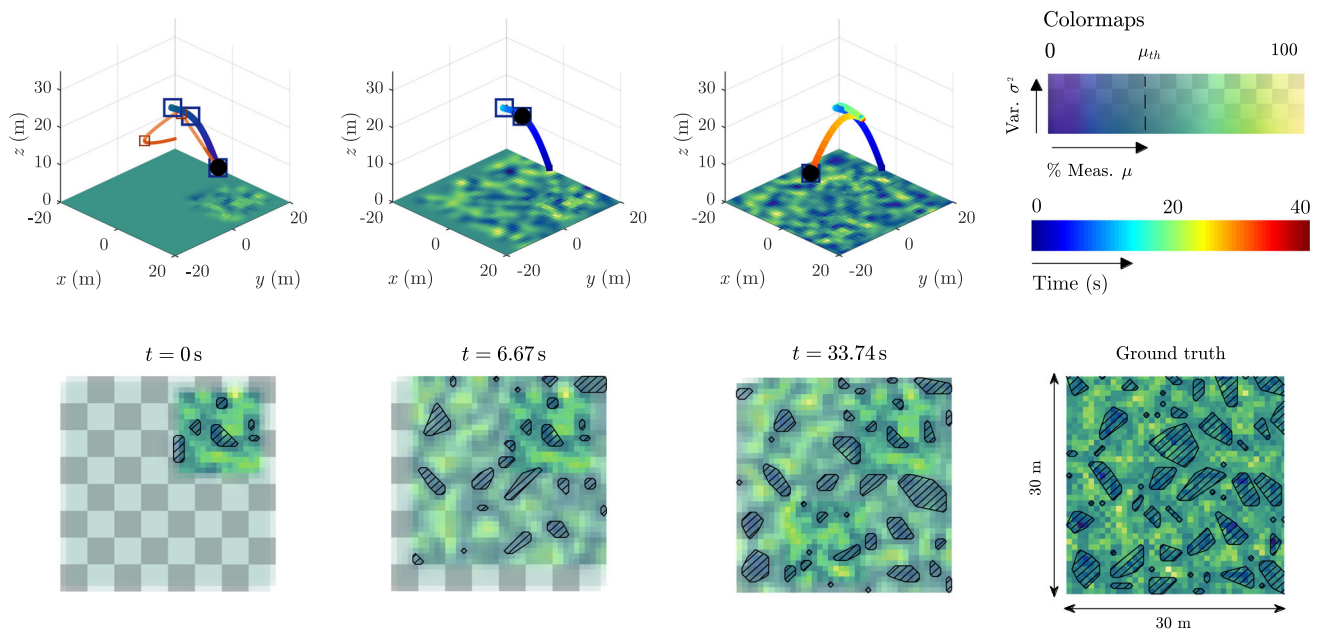


Fig. 6 Example simulation results of our IPP framework. The colormaps are shown on the top-right. Bluer and yellower shades represent lower and higher values of the target parameter, respectively. In the bottom maps, opacity indicates the model uncertainty (variance, σ_i^2), with the checkerboard added for visual clarity and the hatched sections denoting uninteresting areas with $\mu < 40\%$. The ground truth is shown on the bottom-right. The three columns on the left depict the trajectories

(top row) and maps (bottom row) at different snapshots of the mission at times $t = 0$ s, 6.67 s, and 33.74 s. In the top plots, the black dot indicates the current UAV position while the squares show the measurement sites. The top-left figure illustrates an example trajectory before (orange) and after (colored gradient) optimization using the CMA-ES. Note that the map means are rendered in the top trajectory plots (Color figure online)

In the discrete case, this simply amounts to computing Eq. (8) only for the upper set of interesting cells $\mathcal{X}_I = \{\mathbf{x}_i \in \mathcal{X} \mid p_{th} < p(\mathbf{x}_i)\}$, where p_{th} is a threshold on probabilistic occupancy. However, in the continuous case, to account for model uncertainty when planning with the GP model (Eq. (10)), we adopt the principles of bounded uncertainty-aware classification from (Gotovos et al. 2013; Srinivas et al. 2012). The subset of interesting locations \mathcal{X}_I is defined based on the mean and variance of each cell (μ_i, σ_i) as:

$$\mathcal{X}_I = \{\mathbf{x}_i \mid \mathbf{x}_i \in \mathcal{X} \wedge \mu_i + \beta\sigma_i \geq \mu_{th}\}, \quad (11)$$

where μ_{th} is a threshold on the underlying scalar field and β is a design parameter tuned to scale the confidence interval for classification, i.e., the certainty below μ_{th} a cell must possess before being considered interesting.

6 Experimental results

This section discusses our experimental findings in both continuous (Sect. 6.1) and discrete (Sect. 6.2) mapping scenarios. The overall aim is to assess the performance of our framework and demonstrate its flexibility to cater for both types of scenario in missions with and without requirements to adapt

to targeted regions of interest. First, in Sect. 6.1, we evaluate the performance of the proposed approach in simulation and examine the influence of its key parameters. For these experiments, we consider the more complex case of mapping a continuous target variable in a bounded environment using a UAV equipped with an image-based classifier. Sections 6.1.1 and 6.1.2 compare our approach to state-of-the-art methods and study the effects of using different optimization routines in our algorithm. The adaptive replanning scheme is evaluated in Sect. 6.1.3. Then, in Sect. 6.2, we demonstrate the application of our framework for a realistic active classification problem using the RIT-18 dataset (Kemker et al. 2018), before validating real-time capabilities in a real world agricultural monitoring scenario.

To begin, Fig. 6 presents an illustrative example of the progression of our framework for mapping an a priori unknown environment. For adaptive planning, we set a base threshold $\mu_{th} = 40\%$ in Eq. (11) to focus on the more interesting, higher-valued target parameter range. This value allows us to also include unobserved cells in the objective, which are initialized uniformly with a uninformed mean prior of 50%. The top and bottom rows visualize the planned UAV trajectories and maps, respectively, as images are registered at different times during the mission. The top-left plot depicts the first planned trajectory before (orange) and after (colored

gradient) applying the CMA-ES. As shown, the optimization step shifts initial measurement sites (squares) to high altitudes, allowing low-resolution, high-uncertainty data to be quickly collected before the map is refined (second and third columns). A qualitative comparison with ground truth (bottom-right) confirms that our method performs well, producing a fairly complete map in a short time period with most uninteresting regions (hatched areas) identified.

6.1 Simulation results

6.1.1 Comparison against benchmarks

Our framework is evaluated by comparison to benchmarks for continuous variable mapping in a simulated environment. The simulations were run in MATLAB on a single desktop with a 1.8 GHz Intel i7 processor and 16 GB of RAM and model a synthetic information gathering problem in a $30\text{m} \times 30\text{m}$ area. The target distributions are generated as 2-D Gaussian random fields with the mapped scalar parameter ranging from 0 to 100% and cluster radii randomly set between 1 m and 3 m. We use a uniform resolution of 0.75 m for both the training \mathcal{X} and predictive \mathcal{X}^* grids, and perform uninformed initialization with a uniform mean prior of 50%. For the GP, an isotropic Matérn 3/2 kernel is applied. It is defined as (Rasmussen and Williams 2006):

$$k_{\text{Mat3}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \left(1 + \frac{\sqrt{3}d}{l} \right) \exp \left(-\frac{\sqrt{3}d}{l} \right), \quad (12)$$

where d is the Euclidean distance between inputs \mathbf{x} and \mathbf{x}' , and l and σ_f^2 are hyperparameters representing the length scale and signal variance, respectively. We train the hyperparameters $\{\sigma_n^2, \sigma_f^2, l\} = \{1.42, 1.82, 3.67\}$ by maximizing the model log marginal likelihood, using four independent maps with variances modified to cover the entire target parameter range during inference.

For fusing new data, measurement noise is simulated based on the camera model in Fig. 3, with a 10 m altitude beyond which images scale by a factor of $s_f = 0.5$. This places a realistic limit on the quality of data that can be obtained from higher altitudes. We consider a square camera footprint with 60° FoV and a 0.15 Hz measurement frequency. For the purposes of these experiments, we assume no actuation or localization noise and that the on-board camera always faces downwards.

Our approach is compared against three different strategies: (a) the sampling-based rapidly exploring information gathering tree (RIG-tree) introduced by Hollinger and Sukhatme (2014), a state-of-the-art IPP method; (b) a traditional fixed-altitude “lawnmower” coverage path; (c) an upward spiral 3-D coverage path; and (d) random waypoint

selection. The random approach is considered as a naïve benchmark that does not incorporate either the structure or the state of the field map for planning, but can be easily implemented in practice. A 200 s budget B is allocated for all strategies. Considering that, to the best of our knowledge, there is no IPP method that procures optimal results when operating in the continuous trajectory space, we assess performance by comparing different information metrics during a mission. We quantify uncertainty with the trace of the map covariance matrix $\text{Tr}(\mathbf{P})$ and study the Root Mean Squared Error (RMSE) and Mean Log Loss (MLL) at points in \mathcal{X} with respect to ground truth as accuracy statistics. As described by Marchant and Ramos (2014), the MLL is a probabilistic confidence measure incorporating the variance of the predictive distribution. Intuitively, all metrics are expected to reduce as data are acquired over time, with steeper declines signifying better performance.

For our planner and methods (a), (b), and (d), the UAV starting position is specified in the corner of the environment as (7.5 m, 7.5 m) within the field and 8.66 m altitude to assert the same initial conditions as for the complete “lawnmower” coverage pattern. For trajectory optimization, the maximum reference velocity and acceleration are 5 m/s and 2 m/s^2 using polynomials of order $k = 12$, and the number of measurements along a path is limited to 10 for computational feasibility. In our planner, we define polynomials with $N = 5$ waypoints and use the lattice in Fig. 5b for the 3-D grid search. In RIG-tree, we associate control waypoints with vertices, and form polynomials by tracing the parents of leaf vertices to the root. For both planners, we consider the utility $I(\cdot)$ in Eq. (10) and set a base threshold of $\mu_{th} = 40\%$ in Eq. (11) above which map regions are considered interesting to define an adaptive planning requirement.

As outlined in our previous papers (Popović et al. 2017a, b), we use a fixed-horizon version of RIG-tree to allow for incremental replanning and adaptivity and obtain a fair comparison to our planner. We set the UAV starting position as the same corner of the environment as described above, with no prior map information used to create the initial plan. The algorithm alternates between tree construction (replanning) and plan execution, updating the map with each new set of measurements. For replanning, the tree branch expansion step size is set to 10 m with 500 sampling iterations. The latter value was set to obtain the same ~ 20 s replanning time as required by the CMA-ES to optimize a single trajectory, given 45 iterations and the optimization parameters suggested by Hansen (2006). This enables us to compare the methods fairly in the experimental setup based on the time taken to find an informative trajectory.

In the fixed-altitude coverage planner, height (8.66 m) and velocity (0.78 m/s) are defined for complete uniform coverage given the specified budget and measurement frequency. To design a fair benchmark, we studied possible “lawn-

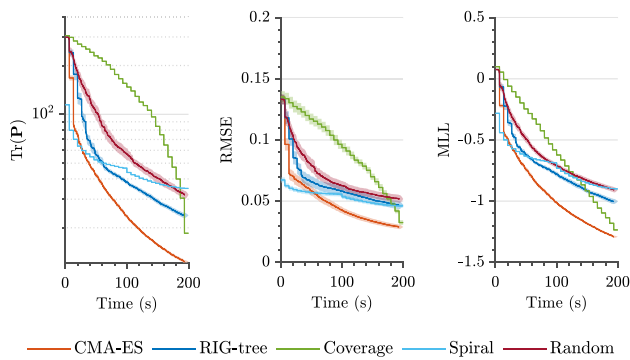


Fig. 7 Comparison of our IPP framework using the CMA-ES against benchmarks for a fixed mission time budget of 200s. The solid lines represent means over 30 trials. The thin shaded regions depict 95% confidence bounds. Using IPP, map uncertainty (left) and error (middle, right) reduce quickly as the UAV obtains low-resolution images before descending. Note the logarithmic scale of the $\text{Tr}(\mathbf{P})$ axis

mower” patterns with heights determined by the camera FoV. For each pattern, we modified velocity to match the budget, then selected the best-performing one. In the 3-D coverage planner, the path is a conical spiral trajectory, reducing in radius with height and spanning the minimum (1 m) and maximum (26 m) UAV flight altitudes. Finally, in the random planner, we randomly sample a destination in the bounded volume above the terrain and generate a trajectory by connecting it to the current UAV position.

Figure 7 shows how the metrics evolve for each planner during the mission. Note that the spiral curve (light blue) is offset at $t = 0$ since the location of the initial measurement is different. For our algorithm, we use the CMA-ES optimization method. The “lawnmower” coverage curve (green) validates our previous results that uncertainty (left) reduces uniformly and deterministically for a constant altitude and velocity. This motivates approaches that permit the UAV to fly at variable altitudes, as they can compromise between sensor uncertainty and FoV. By starting at the lowest altitude before gradually ascending, the spiral planner (light blue) improves the map most quickly early on in the missions but levels off with the lower-quality measurements obtained as the UAV flies further up. In contrast, with a fixed-size camera footprint, the coverage performs better at later stages, but is not capable of achieving the initial rapid map improvement. As expected, both our algorithm (light orange) and RIG-tree (dark blue) perform better than the spiral and random (dark red) benchmarks, given that the latter do not exploit IPP objectives to guide selecting next waypoint destinations.

Our algorithm produces maps with lower uncertainty and error than those of RIG-tree given the same budget. This confirms that our two-stage planner is more effective than sampling-based methods with the proposed mapping strategy. We noted that fixed step size is a key drawback of

RIG-tree, because values allowing initial ascents tend to limit incremental navigation when later refining the map.

Using the same simulation setup, we also conducted a detailed comparison between our approach and “lawnmower” coverage to examine the benefits of IPP for missions of different durations. First, we considered six budgets B (100 s, 200 s, ..., 600 s) on mission time. For each budget, the proposed CMA-ES-based framework was tested over 10 trials, giving a total of 60 simulations, and the coverage planner was run once with its deterministic path. As detailed above, for a fair evaluation, the fixed coverage altitude for each mission time was chosen for best performance among different complete “lawnmower” patterns. As an example, the left plots in Fig. 8 depict the trajectories executed on a 200 s scenario used for the evaluation. The middle graph shows a quantitative analysis of the final achieved map uncertainties ($\text{Tr}(\mathbf{P})$). For comparison, the results of our approach are normalized with the corresponding coverage planner value, so that percentages below 100% (orange line) indicate a better performance of our method. Second, on the right, we compared the mission times needed by the two methods to produce maps with the same final uncertainties. In these experiments, we examined the six fixed budgets for our method (orange) and, for each, investigated the time requirement for the coverage strategy (blue) to obtain the same reconstruction quality.

Figure 8 illustrates two key benefits of using our approach: (a) we obtain maps with lower final uncertainty (middle) by not fixing the flight altitude of the UAV, and (b) we attain significant time savings (right) by allowing for the early collection of low-quality data, as evidenced in Fig. 7. As a result, with increasing mission time, the marginal discrepancy in the final maps increases. The plot on the right shows that our approach also requires substantially less time (> 50% savings for > 500 s missions) to achieve maps with the same uncertainty. This is because the “zig-zag” pattern for these missions must be set at a lower altitude (6.5 m, compared to 8.66 m for lower budgets) to obtain a reduced level of sensor noise, which increases the total distance traveled by the UAV. Interestingly, the coverage path produces a better result only at the end of a 100 s mission; this is because, at this budget, our planner lacks time to descend to refine the map. In future studies, we intend to extend our ideas from previous work (Popović et al. 2017a) and include an awareness of the remaining path budget for planning.

6.1.2 Comparison of optimization methods

Next, we consider the effects of using different optimization routines in Sect. 5.4 to evaluate our CMA-ES-based approach. We use the same simulation setup as described above, i.e., 30 trials for mapping a 30 m × 30 m environment, and study the following methods:

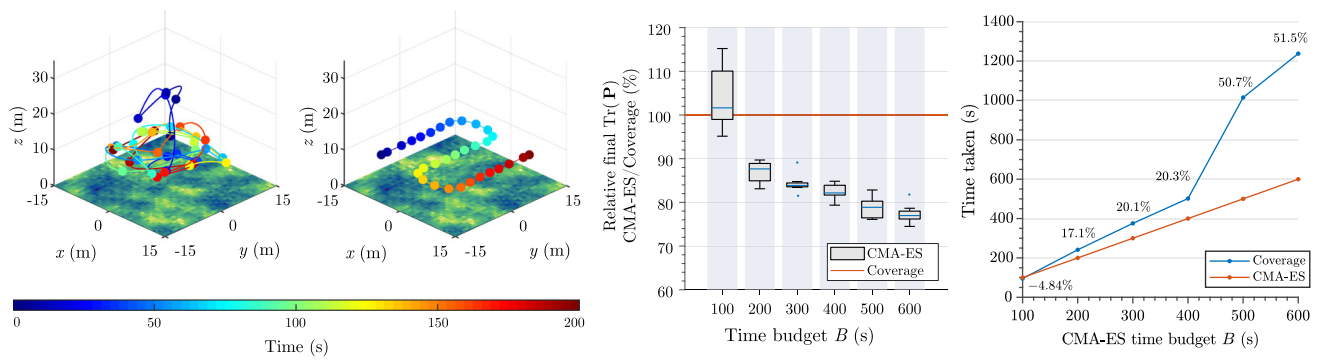


Fig. 8 Left: Example of an evaluation scenario comparing our IPP approach with the CMA-ES to “lawnmower” coverage (left and right plots, respectively) for mapping a continuous variable in 200s missions. The colored lines represent the traveled trajectories, the spheres indicate measurement sites, and the ground truth maps are rendered. Middle: Comparison of the final map uncertainties ($\text{Tr}(\mathbf{P})$) for various path budgets. Ten CMA-ES trials were run for each budget. Right: Com-

parison of times taken to achieve the same final map uncertainty, given a fixed path budget for the CMA-ES. The orange line corresponds to average values over 10 CMA-ES trials, and relative time savings using the new method are shown as percentages. By allowing for altitude variations, our approach trades off between FoV and sensor noise to quickly obtain high-confidence maps with finer end quality in the same time period (Color figure online)

- Lattice*: 3-D grid search only (i.e., without Line 7 in Algorithm 1);
- CMA-ES*: global evolutionary optimization routine (Hansen 2006) (described in Sect. 5.4);
- Interior Point (IP)*: approximate gradient-based optimization using the interior-point approach (Byrd et al. 2006);
- Simulated Annealing (SA)*: global optimization based on the physical cooling process in metallurgy (Ingber and Rosen 1992);
- Bayesian Optimization (BO)*: global optimization using a GP prior (Gelbart et al. 2014).

For optimizers (b)–(e) we investigate initializations using: (i) the lattice search output from (a) and (ii) random point selection in the workspace, in order to investigate their sensitivity to the starting conditions.

The aim of these experiments is to examine how the methods compare using standard MATLAB implementations as baselines. As described above, we set the stopping criteria of the algorithms such that each one is allowed the same ~ 20 s for replanning. Note that the benchmarks were applied using default or recommended values, without significant effort invested into adjusting their parameters, in order to make them practically comparable with the CMA-ES, which requires minimal tuning procedures.

For the local IP optimizer, we approximate Hessians by a dense quasi-Newton strategy and apply the step-wise algorithm described by Byrd et al. (2006). For SA, we apply an exponential cooling schedule and an initial temperature of 100. For BO, we use the time-weighted Expected Improvement acquisition function studied by Gelbart et al. (2014) with an exploration ratio of 0.5. Additionally, for

our approach, we examine two variations of the CMA-ES with initial step sizes of (0.5 m, 0.5 m, 1 m), (3 m, 3 m, 4 m) and (10 m, 10 m, 12 m) in the (x, y, z) co-ordinates, where the z-axis defines altitude. These values were chosen based on the extent of the robot workspace in order to compare different global search behaviors, as the step size parameter effectively captures the distribution from which new solutions are sampled, and thus how well the problem domain is covered by the optimization routine. The aim is to also obtain a practical insight into the tuning requirements for the CMA-ES to achieve best results.

Table 1 displays the mean results for each method averaged over the 30 trials, with the benchmarks from Sect. 6.1.1 included for reference. The suffixes are used to denote the initialization strategy (lattice or random) and different step sizes for the CMA-ES ((0.5, 1), (3, 4), and (10, 12) correspond to step sizes of (0.5 m, 0.5 m, 1 m), (3 m, 3 m, 4 m) and (10 m, 10 m, 12 m), respectively). Following Marchant and Ramos (2014), we also show weighted statistics to emphasize errors in high-valued regions. As the same objective is used for all methods, consistent trends are observed in both non-weighted and weighted metrics. In Fig. 9, we show the mean times taken by each method using the lattice initialization to reduce $\text{Tr}(\mathbf{P})$ to 75% of its initial value to represent the decay speed of the map uncertainty.

Comparing the lattice approach with the CMA-ES and IP methods using this informed initialization (‘lat’) confirms that optimization reduces both uncertainty and error. With the lowest values, the CMA-ES performs best on all indicators as it searches globally to escape local minima. Whereas the CMA-ES variants with smaller and intermediate step sizes behave similarly in this problem, using much larger steps in ‘CMA-ES (10, 12)’ ($> 33\%$ of the workspace extent co-

Table 1 Mean information metrics for all optimization methods, averaged over 30 continuous mapping trials. The suffixes ‘lat.’ and ‘rand.’ indicate initializations using the lattice and random approaches, respectively. The lowest uncertainties and errors obtained with the CMA-ES justify our proposed global optimization strategy

Method	Tr(P)	RMSE	WRMSE	MLL	WMLL
Lattice	51.421	0.0560	0.0559	−0.960	−0.960
CMA-ES (0.5,1) lat.	45.670	0.0520	0.0517	−1.007	−1.009
CMA-ES (3,4) lat.	45.187	0.0525	0.0522	−1.002	−1.006
CMA-ES (10,12) lat.	51.246	0.0603	0.0599	−0.894	−0.897
IP lat.	47.973	0.0541	0.0538	−0.961	−0.962
SA lat.	51.968	0.0571	0.0569	−0.903	−0.904
BO lat.	68.806	0.0685	0.0682	−0.747	−0.749
Random	92.618	0.0771	0.0764	−0.668	−0.669
CMA-ES (3,4) rand.	53.448	0.0613	0.0608	−0.8614	−0.8653
IP rand.	86.709	0.0753	0.0746	−0.658	−0.658
SA rand.	67.730	0.0676	0.0672	−0.749	−0.752
BO rand.	74.581	0.0711	0.0705	−0.693	−0.696
RIG-tree	69.325	0.0690	0.0689	−0.757	−0.757
Coverage	170.911	0.0981	0.0968	−0.665	−0.667
Spiral	54.251	0.0595	0.0595	−0.793	−0.791

Bold indicates the best values (obtained using our method)




Fig. 9 Mean times required by each optimization method to reduce Tr(**P**) (map uncertainty) to 75% of its initial value, averaged over the 30 trials. The CMA-ES, IP, SA, and BO routines use a lattice-based initialization strategy. The CMA-ES result corresponds to a search with intermediate step sizes (‘CMA-ES (3, 4)’). With a mean time of 18.1 s, this approach performs best

ordinate-wise) results in worse performance, as these lead to large random fluctuations during the evolutionary search, which slows down convergence. This reflects the importance of selecting suitable step sizes to cover the application domain, as discussed by Hitz et al. (2017) and Hansen (2006). Surprisingly, applying BO with lattice-based initialization yields mean metrics poorer than those of the lattice itself. We suspect this to be due to its high exploratory behavior causing erratic paths similar to those of ‘CMA-ES (10, 12)’. Despite attempting several commonly used acquisition functions, we found BO to be the most difficult one to tune for the highly nonlinear problem domain.

The fact that the CMA-ES remains as the most successful optimizer using random initialization (‘rand’) suggests that it is most robust to varying starting conditions. Nonetheless, the results show that the lattice search still contributes significantly to finding a good initial solution in practice, which

underlines the benefits of our two-step approach. Though it performs well with the lattice, the IP method using random initialization scores almost as poorly as the random benchmark alone since it only conducts a local search to refine the solution. This implies that it is very dependent on the initial grid selection within our planner, and further motivates global optimization routines.

6.1.3 Adaptive replanning evaluation

We examine the performance of our online framework in planning setups with adaptivity requirements (Sect. 5.5.1) to assess its ability to learn and focus on targeted regions of interest in different environments. The experiments consider two continuous mapping scenarios: (a) ‘Split’, handcrafted maps where the interesting area is well-defined and (b) ‘Gaussian’, the uniformly distributed fields from Sect. 6.1.1. Practically speaking, ‘Gaussian’ maps are representative of cases where the underlying field varies smoothly, matching the assumption of the model, whereas ‘Split’ maps highlight the targeted behaviour of the planner where specific zoning might be present in the field. Both of these scenarios are relevant for an agricultural monitoring application as in Sect. 7, for example, depending on the type of plants growing on a field and the treatments it has received. ‘Split’ maps are partitioned spatially such that half of the cells in \mathcal{X} are classified as interesting based on Eq. (11) with a base threshold of $\mu_{th} = 40\%$ and $\beta = 3$. Then we apply these parameters for adaptive replanning in the simulation setup from Sect. 6.1.1.

The gain of replanning online is evaluated by comparing a *targeted* variant of our approach, i.e., using Eq. (11) as the planning objective, against itself aiming for pure *non-*

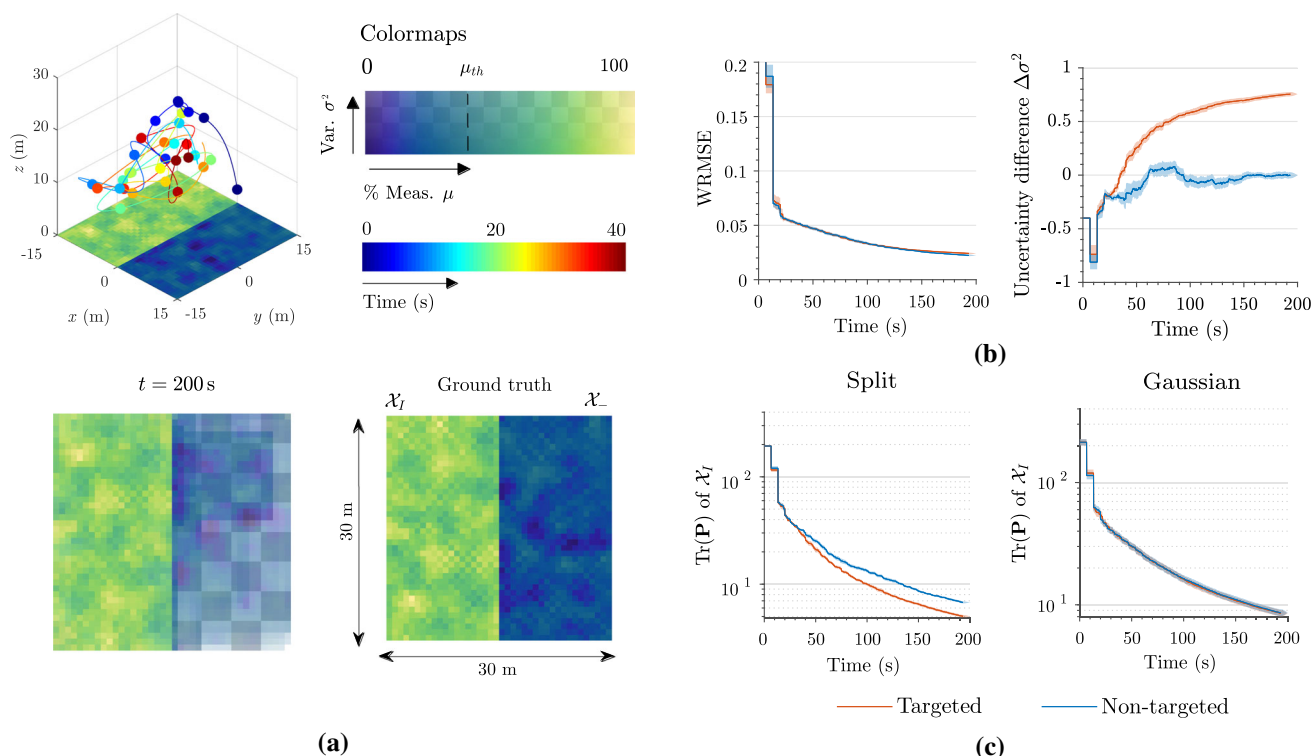


Fig. 10 Evaluation of our adaptive replanning scheme. **a** Top: Example ‘Split’ scenario visualizing the trajectory traveled by our planner using the targeted objective in Eq. (11) (colored line) in a 200 s mission. The spheres indicate measurement sites and the ground truth map is rendered. Colormaps are on the right. The dashed line shows the threshold $\mu_{th} = 40\%$ above which map regions are considered interesting (yellow). **Bottom**: Final map output by our planner (left) compared with ground truth (right). The opacity indicates model uncertainty with the

targeted exploration, i.e., using Eq. (10), which treats the information acquired from all locations in \mathcal{X} equally. As before, we perform 30 simulation trials in $30\text{ m} \times 30\text{ m}$ environments.

For a quantitative evaluation, we consider the variations of WRMSE and the uncertainty difference $\Delta\sigma^2$ in the area of interest and the rest of the total area, which is defined by (Hitz et al. 2017) as:

$$\Delta\sigma^2 = \frac{\bar{\sigma}^2(\mathcal{X}_-) - \bar{\sigma}^2(\mathcal{X}_I)}{\bar{\sigma}^2(\mathcal{X}_-)}, \tag{13}$$

where $\bar{\sigma}^2(\cdot)$ evaluates the mean variance and \mathcal{X}_- and \mathcal{X}_I denote the sets of uninteresting and interesting locations, respectively, as described in Sect. 5.5.1.

Moreover, the rate of uncertainty ($\text{Tr}(\mathbf{P})$) reduction in \mathcal{X}_I evaluates the ability of the planners to focus on interesting regions.

Our results are summarized in Fig. 10. As shown qualitatively in Fig. 10a, once the uninteresting (bluer) map side \mathcal{X}_- is classified in a ‘Split’ environment, planning adaptively in a targeted manner leads to more measurements on

the checkerboard added for visual clarity. Lower opacity confirms higher certainty in the interesting area \mathcal{X}_I . **b** In ‘Split’ scenarios, adaptivity achieves low error (left) with higher uncertainty differences (right) in interesting areas. **c** In ‘Split’ scenarios (left), adaptivity reduces uncertainty faster in interesting areas, while performing comparably to a standard exploration approach in ‘Gaussian’ scenarios (right). In (b) and (c), the solid lines represent means over 30 trials. The shaded regions depict 95% confidence bounds (Color figure online)

the interesting (yellow) side \mathcal{X}_I , which induces lower final uncertainty in this area, as expected. Figure 10b confirms that, in these scenarios, the relative uncertainty difference $\Delta\sigma^2$ increases more rapidly using adaptivity, while map WRMSE remains low. Note that, early in the mission ($< 30\text{ s}$), both approaches behave similarly to explore the initially unknown map. Finally, Fig. 10c shows that the benefit of adaptivity, in terms of reducing uncertainty in areas of interest, is higher in ‘Split’ environments when compared with ‘Gaussian’. Since the region \mathcal{X}_I is clearly distinguished, purely informative measurements can be taken within the camera FoV given the thresholded objective. Planning adaptively, however, yields no disadvantages when the field is uniformly dispersed.

6.2 RIT-18 mapping scenario

We demonstrate our complete framework on a photorealistic scenario in the Gazebo-based RotorS simulation environment (Furrer et al. 2016). In contrast to the preceding section, these experiments show our framework in the discrete mapping domain to reflect the nature of the target dataset.




Fig. 11 Our photorealistic simulation setup in RotorS. **a**, **b** depict the AscTec Firefly UAV and the view from its on-board camera. **c** Shows an aerial view of the $200\text{m} \times 290\text{m}$ surveyed area (RIT-18 validation orthomosaic). The red and green lines annotate the two target classes for mapping using our approach: ‘Lake’ and ‘BRV’, respectively (Color figure online)

Figure 11 depicts our experimental setup, which runs on a single desktop with a 2.6GHz Intel i7 processor and 16GB of RAM. The planning and mapping algorithms were implemented in MATLAB on Ubuntu Linux and interfaced to the Robot Operating System. For mapping, we use RIT-18 (Kemker et al. 2018), a high-resolution 6-band VNIR dataset for semantic segmentation consisting of coastal imagery along Lake Ontario in Hamlin, NY. In our simulations, the surveyed region is a $200\text{m} \times 290\text{m}$ area featuring the RIT-18 validation fold.

Our UAV model is an AscTec Firefly equipped with a downward-facing camera, which has a $360\text{px} \times 480\text{px}$ image resolution and a $(35.4^\circ, 47.2^\circ)$ FoV in the x - and y -directions, respectively. To extract measurements for active classification, we use a modified version of the SegNet convolutional architecture (Badrinarayanan et al. 2017; Sa et al. 2018) accepting multispectral as well as RGB image inputs. The imagery registered from a given UAV pose is passed to the network to produce a dense semantic segmentation output, as exemplified in Fig. 12. We simplify the classification problem by only mapping the following 3 classes derived from the RIT-18 labels: (a) ‘Lake’; (b) a combination of ‘Building’, ‘Road Markings’, and ‘Vehicle’ (‘BRV’); and (c) ‘Background’ (‘Bg’), i.e., all others. These particular labels were chosen based on their distributions to obtain strongly altitude-dependent classification performance, as relevant for the terrain monitoring problem setup.

To model the sensor for predictive planning, we first trained SegNet on all labels using RIT-18 training fold imagery. Our training procedure uses (323, 72, 16) images simulated at 3 different altitudes, (50 m, 70 m, 100 m), and is




Fig. 12 Example classification result from an altitude of 70 m. **a** Shows the RGB image channel input, and **b** visualizes the dense segmentation output. In **(b)**, the probabilistic output for each class [‘Lake’, ‘BRV’, ‘Bg’] is mapped to the corresponding pixel intensity on the [R, G, B] channels




Fig. 13 Sensor models for our trained classifier for the **a** ‘Lake’ and **b** ‘BRV’ classes. The blue and orange curves depict the probability of observing class label ‘1’, given that the map contains ‘1’ or ‘0’, i.e., true and false positives, respectively. Note that the false positive probability can decrease with altitude as the classifier becomes more conservative with true outputs (Color figure online)

performed on an Nvidia Titan X Pascal GPU module. At 70 m and 100 m, the training images were additionally downsampled to exaggerate the effects of pixel mixing at lower resolutions. Then, classification accuracy was assessed by using validation data to compute confusion matrices at each altitude for the 3 classes of interest (30% train and 70% test split, with a higher proportion of training data at lower altitudes). This enabled us to derive the sensor models in Fig. 13, in which we associate intermediate altitudes with the closest performance statistics available. Note that the altitude range considered here is wider compared to the previous sections due to the larger environment size.

We employ a discrete strategy to map the target region, maintaining one independent occupancy grid layer for each of the 3 classes. Each layer has a uniform resolution of 5 m, and all cells are initialized with an uninformed probability of 0.5. For predictive measurements when planning, we use the sensor models in Fig. 13 conditioned on the most probable current map states. For fusing new data, we project the classifier output in Fig. 12b on the occupancy grids for each class, performing likelihood updates with the maximum pixel probabilities mapping to each cell. Note that, unlike the pixel-

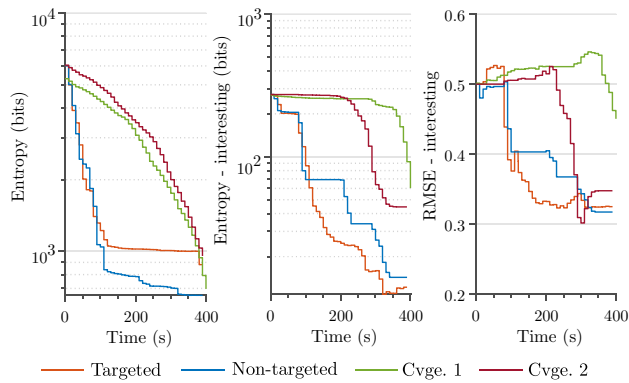


Fig. 14 Comparison of our IPP approach using the CMA-ES against fixed-altitude coverage benchmarks ('Cvge. 1' = 157 m, 'Cvge. 2' = 104 m) in a 400 s photorealistic mapping scenario. By planning adaptively, map uncertainty (middle) and error (right) in interesting areas ('BRV' class) reduce most rapidly, while yielding higher overall map uncertainty (left). Note the logarithmic scale of the Entropy axis

wise classifier output, our mapping strategy does not enforce the probabilities of a cell across the layers to sum to 1, as a cell may contain objects from multiple classes.

The planning goal in this setup is to efficiently map the 'BRV' class, which would be useful, e.g., for identifying man-made features in search and rescue scenarios. Our proposed approach with the CMA-ES is evaluated against the "lawnmower" coverage strategy, considered as the naïve choice of algorithm for such applications. To investigate height-dependent performance, trials are performed with two coverage patterns at fixed altitudes of 157 m and 104 m, denoted 'Cvge. 1' and 'Cvge. 2', respectively. In addition, we study both targeted and non-targeted versions of our approach, in order to expose the benefits of using adaptive replanning to map the regions of interest. Our performance metrics are map entropy and RMSE with respect to the RIT-18 ground truth labels.

All methods are given a 400 s budget B . To limit computational load on the classifier, we assign a measurement frequency of 0.1 Hz, allowing the UAV to stop while processing images. As before, trajectory optimization is performed on polynomials of order $k = 12$. The UAV starting position in our approach is set as (33 m, 46 m) within the lower-left field corner with 104 m altitude to achieve consistency with the lower-altitude coverage pattern. For planning, we use polynomials defined by $N = 5$ waypoints with a reference velocity and acceleration of 15 m/s and 20 m/s². The 3-D grid search is executed on a scaled version of the 30-point lattice in Fig. 5b, stretched to cover the rectangular area, and the CMA-ES optimizer runs with initial step sizes of (50 m, 60 m, 40 m). We apply a low threshold of $p_{th} = 0.4$ in Eq. (9) on the occupancy grid layer of the target 'BRV' class to define the adaptive planning requirement. The coverage benchmarks are designed based on the principles discussed in the preceding sections.

Figure 14 compares the performance of each planner in this scenario. As in Sect. 6.1.1, total map uncertainty reduces uniformly using the coverage patterns, with interesting areas surveyed only towards the end due to the environment layout. In these regions, 'Cvge. 2' (dark red) achieves higher-quality mapping than 'Cvge. 1' (green) as its lower altitude permits more accurate measurements. This evidences the height-dependent nature of the classifier, which motivates using IPP to navigate in 3-D space. By planning adaptively using our targeted approach (orange), both uncertainty and error decay most rapidly in the areas of interest, as expected. However, a non-targeted strategy (blue) performs better in terms of overall map uncertainty, since it is biased towards pure exploration. These results demonstrate how our framework can be tailored to balance exploration (uniform uncertainty reduction) and exploitation (mapping a target class) in a particular scenario.

Figure 15 visualizes the trajectory traveled by our targeted planner during the mission. As before, the UAV initially (< 100 s) explores unobserved space, before concentrating on high-probability areas for the 'BRV' class once they have been discovered (green, or cyan for cells containing both 'BRV' and 'Bg'). It can be seen that the map becomes more complete in these regions as low-altitude measurements are accumulated. Note that the two small cars to the right of the building and above the parking lot (visible in Fig. 11c) are mapped incorrectly as our SegNet model is limited in segmenting out fine details given the data it was trained on. Considering the richness of the RIT-18 dataset, an interesting direction for future work is to explore different classification methods and target classes within our IPP framework.

7 Field deployment

Finally, we present experimental results from a field deployment implementing our framework on a UAV to monitor the vegetation distribution in a field. The aim is to validate the system for performing a practical sensing task in a challenging outdoor environment, with all algorithms running on-board and in real-time.

The experiments were conducted on an agricultural field at the Research Station for Plant Sciences Lindau of ETH Zurich in Switzerland (Lat. 47.450040°, Lon. 8.681056°) in Sept. 2018. Figure 16 depicts our on-field setup. The UAV platform in (a) is the DJI Matrice M100 monitoring a 20 m × 20 m area within the field with maximum and minimum altitudes of 21 m and 8 m. As shown in (b), the controlled field features a central area of crops planted in row arrangements, surrounded by dense weed distributions on its edges. The width of the crop row area is ~ 18 m. Vegetation mapping is performed using RGB imagery from a downward-facing Intel RealSense ZR300 camera with a res-

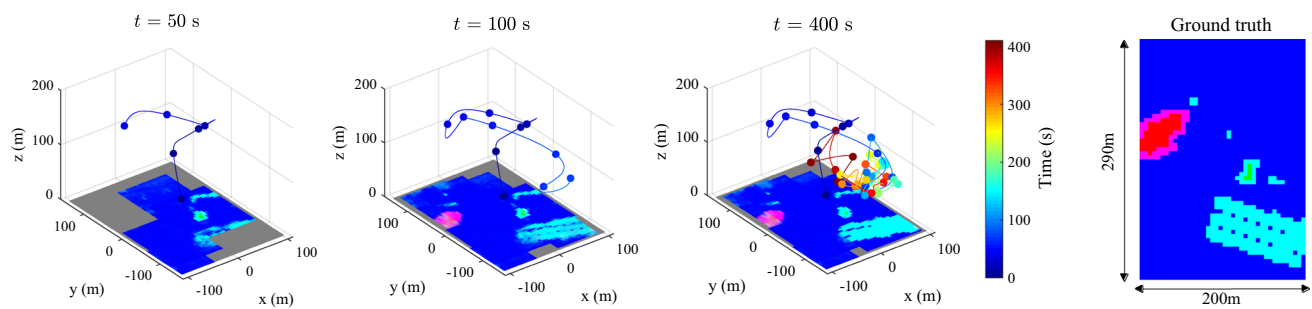


Fig. 15 Visualization of the trajectory traveled by our targeted planner (colored line) in the 400 s mission. The three plots depict different snapshots of the mission at times $t = 50$ s, 100 s, and 400 s. The RIT-18 ground truth is shown on the right. In the trajectory plots, the spheres indicate measurement sites and the current occupancy map states are rendered. The colors portray composites of the three-layer map representation, with the cell probabilities for each class ['Lake' , 'BRV' ,

'Bg'] mapped to the corresponding intensities on the [R, G, B] channels. Gray indicates unobserved space. The sequence shows that our planner quickly explores the area to later focus on more closely mapping interesting regions ('BRV' class). Note that the magenta and cyan cells indicate the presence of two classes, ['Bg' , 'Lake'] and ['Bg' , 'BRV'], respectively (Color figure online)




Fig. 16 **a** shows a close-up of the DJI Matrice M100 used in the outdoor trials. **b** illustrates the UAV executing a mission, flying to monitor the vegetation distribution in a field of interest. **c**, **d** exemplify RGB images taken by the on-board camera at different altitudes, which are used for mapping online. Note that, in (**b**), the UAV is positioned approximately above a boundary between the controlled crop (left) and weed (right) areas, visible in (**c**, **d**)

olution of $1920\text{px} \times 1080\text{px}$ and a FoV of $(68.0^\circ, 47.2^\circ)$. Example images taken from different altitudes are shown in (**c**) and (**d**). In these images, the boundaries between the weed and crop areas of the field can be distinguished.

We use the robust visual inertial odometry (ROVIO) framework (Bloesch et al. 2015) for state estimation with model predictive control (MPC) (Kamel et al. 2017) to track trajectories output by the planner. All computations, including modules for environmental mapping and informative planning, are based on our open-source package and run on an on-board computer with a 3.2GHz Intel NUC i7, 16GB of RAM, and running Ubuntu Linux 16.04 LTS with Robot Operating System (ROS) as middleware. Further platform specifics are discussed by Sa et al. (2017, 2018).

The goal is to map the level of Excess Green Index (ExG) in the area of interest based on the RGB images and using our GP-based mapping method for continuous variables. It is defined by $ExG = 2g - r - b$, where r , g , and b are the normalized red, green, and blue color channels in the RGB space (Yang et al. 2015). Note that we normalized the range of the ExG based on the maximum and minimum magnitudes measured on the experimental field in previously acquired datasets.

For mapping, a uniform resolution of 0.5 m is set for both the training \mathcal{X} and predictive \mathcal{X}^* grids in the GP. Assuming that the field only contains soil, we initialize the map with a uniform mean prior of 0 normalized ExG. The isotropic Matérn 3/2 kernel (Eq. (12)) is applied to approximate the vegetation spread in the field, with the hyperparameters $\{\sigma_n^2, \sigma_f^2, l\} = \{0.50, 0.50, 1.76\}$ trained using images from a manually acquired dataset. We recorded the dataset while flying the UAV at a fixed altitude of 8 m over the area, which corresponds to the minimum level allowed in the experiments, and the maximum GSD (in m/px) obtainable in the images, i.e., highest possible mapping resolution in \mathcal{X} .

Measurements are extracted from RGB images taken at a constant frequency of 0.20 Hz. We consider the sensor noise model defined by Eq. (7) with coefficients $a = 0.05$ and $b = 0.2$ set to determine the variance variation based on the altitude range. Note that, due to the inherently high camera resolution, the scaling factor is $s_f = 1$ across the altitude range. To update the map with a new image, the pixels are first projected on the ground plane based on the camera parameters and the estimated UAV pose. The projected pixels are averaged per cell, and the normalized ExG values for each cell are computed from the color channels. Data fusion then proceeds according to the method described in Sect. 4.2.2.

The planning strategy considers polynomial trajectories of order $k = 12$ defined by $N = 3$ waypoints and optimized for a maximum reference velocity and acceleration of 5 m/s and

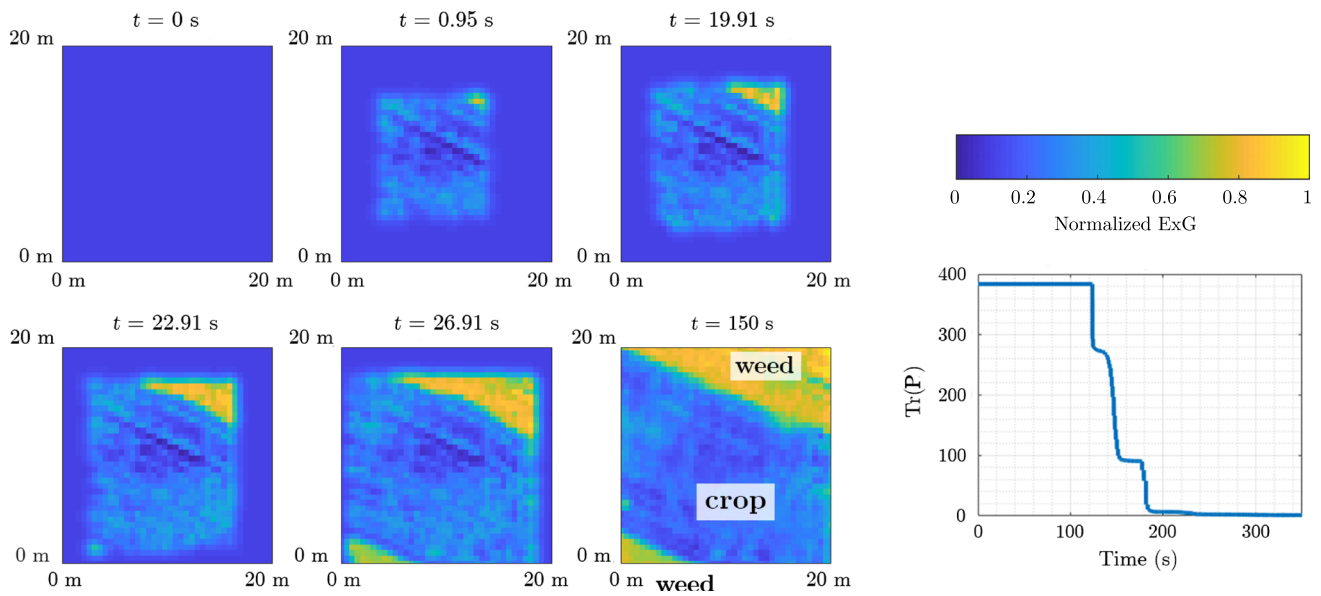


Fig. 17 Experimental results from the field trial using our IPP framework for UAV-based vegetation monitoring. Top-right: Colorbar, indicating the normalized excess green index (ExG) level (plant greenness). Left: Estimated map reconstructions (GP means) of the normalized ExG on the field at different times t during the mission. The bluer and yellow

lower regions correspond to successfully identified crop row and weed regions, respectively. Bottom-right: Evolution of total map uncertainty over time. The results confirm improving map completeness as the number of measurements increases (Color figure online)

3 m/s^2 . The planning objective is set as uncertainty reduction with no interest-based threshold (Eq. (10)), i.e., the aim is to reconstruct the field in a uniform manner, as quickly as possible. We perform the 3-D grid search over a coarse 14-point lattice and set initial CMA-ES step sizes of 4.5 m in each co-ordinate of the UAV workspace. For simplicity, we do not specify a mission budget B ; instead allowing the algorithm to create fixed-horizon plans until the mapping output is perceived visually as being complete.

The deployment results are reported in Fig. 17². Since the ground truth data of normalized ExG levels are not available, we validate our framework by assessing the progression of total map uncertainty on the bottom-right, which confirms that uncertainty is reduced over time. Note that the curve in the plot is offset by ~ 100 s as data recording was triggered before the UAV took off to take the first measurement.

Qualitatively, the sequence of plots on the left verifies that the estimated map does become more complete as images of the field are accumulated. The yellower parts, corresponding to areas with high values of the normalized ExG, indicate successfully identified weeds on the edges of the physical field (visible in Fig 16c, d). Towards the central area, a close look at the bluer parts with lower ExG reveals that even the crop row details are mapped correctly. These findings showcase our mapping strategy and represent valuable data that could guide practical crop management decisions.

² A video showing the UAV trajectory is available at: <http://youtu.be/5dK8LcQH85o>.

8 Conclusion and future work

This paper introduced a general IPP framework for environmental monitoring applications using an aerial robot. The method is capable of mapping either discrete or continuous target variables on a terrain using variable-resolution data received from probabilistic sensors. The resulting maps are employed for IPP by optimizing parameterized continuous-space trajectories initialized by a coarse 3-D search.

Our approach was evaluated extensively in simulations using synthetic and real world data. The results reveal higher efficiency compared to state-of-the-art methods and highlight its ability to efficiently build models with lower uncertainty in value-dependent regions of interest. Furthermore, we validated our framework in an active classification problem using a publicly available dataset. These experiments demonstrated its online application on a photorealistic mapping scenario with a SegNet-based sensor for data acquisition. Finally, a proof of concept was presented showing the algorithms running on-board and in real-time for an agricultural monitoring task in an outdoor environment.

The implementation of the proposed planner is released for use and further development by the community along with sample experimental results. Future theoretical work will investigate scaling the approach to larger environments and extending the mapping model to capture temporal dynamics. This would enable previously acquired data to be used as a prior in persistent monitoring missions. Towards more accurate map building in practice, it would be interesting

to also incorporate the robot localization uncertainty in the decision-making algorithm.

Acknowledgements We would like to thank Dr. Frank Liebisch for his useful discussions. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644227 and from the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0029.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Arora, S., Choudhury, S., & Scherer, S. (2018). Hindsight is only 50/50: Unsuitability of MDP based approximate POMDP solvers for multi-resolution information gathering. [arXiv:1804.02573](https://arxiv.org/abs/1804.02573).
- Arora, S., & Scherer, S. (2017). Randomized algorithm for informative path planning with budget constraints. In *IEEE international conference on robotics and automation* (pp. 4997–5004). IEEE.
- Badrinarayanan, V., Alex, K., & Cipolla, R. (2017). SegNet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495.
- Bellini, A. C., Lu, W., Naldi, R., & Ferrari, S. (2014). Information driven path planning and control for collaborative aerial robotic sensors using artificial potential functions. In *American control conference* (pp. 590–597).
- Berrio, J. S., Ward, J., Worrall, S., Zhou, W., & Nebot, E. (2017). Fusing lidar and semantic image information in octree maps. In *Australasian conference on robotics and automation*. ACRA.
- Binney, J., & Sukhatme, G. S. (2012). Branch and bound for informative path planning. In *IEEE international conference on robotics and automation* (pp. 2147–2154). IEEE.
- Bircher, A., Siegwart, R., Kamel, M., Alexis, K., & Oleynikova, H. (2016). Receding horizon path planning for 3D exploration and surface inspection. *Autonomous Robots*, 42(2), 291–306.
- Bloesch, M., Omari, S., Hutter, M., & Siegwart, R. (2015). Robust visual inertial odometry using a direct EKF-based approach. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 298–304). IEEE.
- Byrd, R. H., Gilbert, J. C., Nocedal, J., Byrd, R. H., Gilbert, J. C., Nocedal, J., et al. (2006). A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming*, 89(1), 149–185.
- Cai, C., & Ferrari, S. (2009). Information-driven sensor path planning by approximate cell decomposition. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(3), 672–689.
- Charrow, B., Liu, S., Kumar, V., & Michael, N. (2015). Information-theoretic mapping using cauchy-schwarz quadratic mutual information. In *IEEE international conference on robotics and automation* (pp. 4791–4798). IEEE.
- Chekuri, C., & Pál, M. (2005). A recursive greedy algorithm for walks in directed graphs. In *IEEE symposium on foundations of computer science* (pp. 245–253). IEEE.
- Chen, M., Frazzoli, E., Hsu, D., & Lee, W. S. (2016). POMDP-lite for robust path planning under uncertainty. In *IEEE international conference on robotics and automation* (pp. 5427–5433). IEEE.
- Choset, H. (2001). Coverage for robotics: a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31, 113–126.
- Colomina, I., & Molina, P. (2014). Unmanned aerial systems for photogrammetry and remote sensing: A review. *ISPRS Journal of Photogrammetry and Remote Sensing*, 92, 79–97.
- Dunbabin, M., & Marques, L. (2012). Robots for environmental monitoring: Significant advancements and applications. *IEEE Robotics and Automation Magazine*, 19(1), 24–39.
- Elfes, A. (1989). Using occupancy grids for mobile robot perception and navigation. *Computer*, 22(6), 46–57.
- Ezequiel, C. A. F., Cua, M., Libatique, N. C., Tangonan, G. L., Alampay, R., Labuguen, R. T., Favila, C. M., Honrado, J. L. E., Canos, V., Devaney, C., Loreto, A. B., Bacusmo, J., & Palma, B. (2014). UAV aerial imaging applications for post-disaster assessment, environmental management and infrastructure development. In *International conference on unmanned aircraft systems* (pp. 274–283). IEEE.
- Ferrari, S., & Cai, C. (2009). Information-driven search strategies in the board game of CLUE. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 39(3), 607–625.
- Furrer, F., Burri, M., Achtelik, M., & Siegwart, R. (2016). RotorS—A modular gazebo MAV simulator framework. In *Robot operating system (ROS): The complete reference* (Vol. 1, pp. 595–625). Springer.
- Galceran, E., & Carreras, M. (2013). A survey on coverage path planning for robotics. *Robotics and Autonomous Systems*, 61(12), 1258–1276.
- Gao, M., Xu, X., Klinger, Y., Van Der Woerd, J., & Tapponnier, P. (2017). High-resolution mapping based on an unmanned aerial vehicle (UAV) to capture paleoseismic offsets along the Altyn-Tagh fault, China. *Scientific Reports*, 7(1), 1–11.
- Gelbart, M. A., Snoek, J., & Adams, R. P. (2014). Bayesian optimization with unknown constraints. In *Conference on Uncertainty in Artificial Intelligence* (pp. 250–259). AUAI Press.
- Girdhar, Y., & Dudek, G. (2015). Modeling curiosity in a mobile robot for long-term autonomous exploration and monitoring. *Autonomous Robots*, 33(4), 645–657.
- Gotovos, A., Casati, N., Hitz, G., & Krause, A. (2013). Active learning for level set estimation. In *International joint conference on artificial intelligence* (pp. 1344–1350). AAAI Press.
- Grocholsky, B., Makarenko, A., & Durrant-Whyte, H. (2003). Information-theoretic coordinated control of multiple sensor platforms. In *IEEE international conference on robotics and automation* (pp. 1521–1526). IEEE.
- Hansen, N. (2006). The CMA evolution strategy: A comparing review. *Studies in Fuzziness and Soft Computing*, 192(2006), 75–102.
- Hansen, N. (2009). A method for handling uncertainty in evolutionary optimization with an application to feedback control of combustion. *IEEE Transactions on Evolutionary Computation*, 13(2), 180–197.
- Hitz, G., Galceran, E., Garneau, M. È., Pomerleau, F., & Siegwart, R. (2017). Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics*, 34(8), 1427–1449.
- Hitz, G., Gotovos, A., Pomerleau, F., Garneau, M. E., Pradalier, C., Krause, A., & Siegwart, R. Y. (2014). Fully autonomous focused exploration for robotic environmental monitoring. In *IEEE international conference on robotics and automation* (pp. 2658–2664). IEEE.

- Hollinger, G. A., Englot, B., Hover, F. S., Mitra, U., & Sukhatme, G. S. (2013). Active planning for underwater inspection and the benefit of adaptivity. *The International Journal of Robotics Research*, 32(1), 3–18.
- Hollinger, G. A., & Sukhatme, G. S. (2014). Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9), 1271–1287.
- Ingber, L., & Rosen, B. (1992). Genetic algorithms and very fast simulated reannealing: A comparison. *Mathematical and Computer Modelling*, 16(11), 87–100.
- Jadidi, M. G., Miro, J. V., & Dissanayake, G. (2016). Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring. [arXiv:1607.01883](https://arxiv.org/abs/1607.01883).
- Jadidi, M. G., Miro, J. V., & Dissanayake, G. (2018). Gaussian process autonomous mapping and exploration for range sensing mobile robots. *Autonomous Robots*, 42(2), 273–290.
- Julian, B. J., Angermann, M., Schwager, M., & Rus, D. (2011). A scalable information theoretic approach to distributed robot coordination. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 5187–5194). IEEE.
- Kaelbling, L., Littman, M., & Cassandra, A. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1–2), 99–134.
- Kamel, M., Stastny, T., Alexis, K., & Siegwart, R. (2017). Model predictive control for trajectory tracking of unmanned aerial vehicles using robot operating system. In *Robot operating system (ROS): The complete reference* (Vol. 2, pp. 3–39). Springer.
- Kemker, R., Salvaggio, C., & Kanan, C. (2018). Algorithms for semantic segmentation of multispectral remote sensing imagery using deep learning. *ISPRS Journal of Photogrammetry and Remote Sensing*, 145, 60–77.
- Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9, 235–284.
- Kurniawati, H., Hsu, D., & Lee, W. S. (2008). SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*. MIT Press. <http://www.roboticsproceedings.org/rss04/p9.html>.
- Lim, Z. W. (2015). Planning under uncertainty: From informative path planning to partially observable semi-MDPs. PhD thesis, National University of Singapore.
- Lim, Z. W., Hsu, D., & Lee, W. S. (2015). Adaptive informative path planning in metric spaces. *The International Journal of Robotics Research*, 35, 585–598.
- Lu, W., Zhang, G., & Ferrari, S. (2014). An information potential approach to integrated sensor path planning and control. *IEEE Transactions on Robotics*, 30(4), 919–934.
- Manfreda, S., McCabe, M. F., Miller, P. E., Lucas, R., Madrigal, V. P., Mallinis, G., et al. (2018). On the use of unmanned aerial systems for environmental monitoring. *Remote Sensing*, 10(4), 641.
- Marchant, R., & Ramos, F. (2014). Bayesian optimisation for informative continuous path planning. In *IEEE international conference on robotics and automation* (pp. 6136–6143). IEEE.
- Morere, P., Marchant, R., & Ramos, F. (2017). Sequential bayesian optimisation as a POMDP for environment monitoring with UAVs. In *IEEE international conference on robotics and automation* (pp. 6381–6388). IEEE.
- O’Callaghan, S. T., & Ramos, F. T. (2012). Gaussian process occupancy maps. *The International Journal of Robotics Research*, 31(1), 42–62.
- Popović, M., Hitz, G., Nieto, J., Sa, I., Siegwart, R., & Galceran, E. (2017a). Online informative path planning for active classification using UAVs. In *IEEE international conference on robotics and automation* (pp. 5753–5758). IEEE.
- Popović, M., Vidal-Calleja, T., Hitz, G., Sa, I., Siegwart, R. Y., & Nieto, J. (2017b). Multiresolution mapping and informative path planning for UAV-based terrain monitoring. In *IEEE/RSJ international conference on intelligent robots and systems* (pp. 1382–1388). IEEE.
- Rasmussen, C. E., & Williams, C. K. I. (2006). *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
- Rastgoftar, H., Zhang, B., & Atkins, E. M. (2018). A data-driven approach for autonomous motion planning and control in off-road driving scenarios. [arXiv:1805.09951](https://arxiv.org/abs/1805.09951).
- Reece, S., & Roberts, S. (2013). An Introduction to Gaussian processes for the Kalman filter expert. In *FUSION* (pp. 1–9).
- Richter, C., Bry, A., & Roy, N. (2013). Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments. In *International symposium of robotics research* (pp. 649–666). Springer.
- Sa, I., Chen, Z., Popović, M., Khanna, R., Liebisch, F., Nieto, J., & Siegwart, R. (2018). weedNet: Dense semantic weed classification using multispectral images and MAV for smart farming. In *IEEE robotics and automation letters* (pp. 588–595). IEEE.
- Sa, I., Kamel, M., Khanna, R., Popović, M., Nieto, J., & Siegwart, R. (2017). Dynamic System Identification, and Control for a cost effective open-source VTOL MAV. In *Field and Service Robotics*. Springer.
- Sadat, S. A., Wawerla, J., & Vaughan, R. (2015). Fractal trajectories for online non-uniform aerial coverage. In *IEEE international conference on robotics and automation* (pp. 2971–2976). IEEE.
- Schwager, M., Dames, P., Rus, D., & Kumar, V. (2017). *A multi-robot control policy for information gathering in the presence of unknown hazards* (pp. 455–472). New York: Springer.
- Sim, R., & Roy, N. (2005). Global A-optimal robot exploration in SLAM. In *IEEE international conference on robotics and automation* (pp. 661–666). IEEE.
- Singh, A., Krause, A., Guestrin, C., & Kaiser, W. J. (2009). Efficient informative sensing using multiple robots. *Journal of Artificial Intelligence Research*, 34, 707–755.
- Singh, A., Ramos, F., Whyte, H. D., & Kaiser, W. J. (2010). Modeling and decision making in spatio-temporal processes for environmental surveillance. In *IEEE international conference on robotics and automation* (pp. 5490–5497). IEEE.
- Srinivas, N., Krause, A., Kakade, S. M., & Seeger, M. W. (2012). Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. In *IEEE transactions on information theory* (Vol. 58, No. 5, pp. 3250–3265). IEEE.
- Vidal-Calleja, T., Su, D., Bruijn, F. D., & Miro, J. V. (2014). Learning spatial correlations for bayesian fusion in pipe thickness mapping. In *IEEE international conference on robotics and automation* (pp. 683–690). IEEE.
- Vivaldini, K. C. T., Guizilini, V., Oliveira, M. D. C., Martinelli, T. H., Wolf, D. F., & Ramos, F. (2016). Route planning for active classification with UAVs. In *IEEE international conference on robotics and automation* (pp. 2563–2568). IEEE.
- Wei, H., Lu, W., Zhu, P., Ferrari, S., Liu, M., Klein, R. H., et al. (2016). Information value in nonparametric Dirichlet-process Gaussian-process (DPGP) mixture models. *Automatica*, 74, 360–368.
- Yamauchi, B. (1998). Frontier-based exploration using multiple robots. In *International conference on autonomous agents (AGENTS ’98)*, pp. 47–53. ACM.
- Yang, W., Wang, S., Zhao, X., Zhang, J., & Feng, J. (2015). Greenness identification based on HSV decision tree. *Information Processing in Agriculture*, 2(3–4), 149–160.
- Zhang, H., & Vorobeychik, Y. (2016). Submodular optimization with routing constraints. In *AAAI Conference on Artificial Intelligence*, (pp. 819–825). AAAI Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Marija Popović is a Research Associate at the Smart Robotics Lab at Imperial College London. She received her Ph.D. from ETH Zurich in 2019 and her Master in Engineering in Integrated Mechanical and Electrical Engineering from the University of Bath in 2015. Her research interests lie in the areas of active perception and path planning for autonomous mobile systems, focusing on applications in aerial robotics.



Teresa Vidal-Calleja received her B.Sc. in Mechanical Engineering from the National Autonomous University of Mexico (UNAM), her MSc in Electrical Engineering (Mechatronics options) from CINVESTAV-IPN, Mexico City, and her Ph.D. in Automatic Control, Computer Vision and Robotics from the Polytechnic University of Catalonia (UPC) in 2007. She joined the Centre for Autonomous Systems (CAS) at University of Technology Sydney (UTS) in 2012, where she was

UTS Chancellor's Research Fellow and later became Senior Lecturer. Her research interests include robotic perception, automatic recognition, visual SLAM, aerial and ground robots cooperation, and autonomous navigation.



Gregory Hitz is the CEO of Sevensense Robotics in Zurich. Before that, he was a senior researcher at the Autonomous Systems Lab in ETH Zurich, where he also received his M.Sc. (2011) and Ph.D. (2016) degrees. His research interests range from state estimation and mapping to path planning, particularly in the informative setting. He is passionate about bringing robots out of the research lab into real applications.



and B.E. (2010) from the University of Sydney.

Jen Jen Chung is a Senior Researcher in the Autonomous Systems Lab at ETH Zurich. Her current research interests include multi-robot coordination methods and risk-aware planning for aerial and ground robots. Her past research includes multiagent learning methods at Oregon State University and information-based exploration-exploitation strategies for autonomous soaring platforms at the Australian Centre for Field Robotics in the University of Sydney. She received her Ph.D. (2014)



solve challenging tasks in real-world environments. A particular focus of his research has been on enabling aerial and ground robots to perform an autonomous survey/navigate in agricultural or underground subterranean scenarios.

Inkyu Sa is a research scientist at the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Robotics and Autonomous Systems Group, Brisbane, Australia, since 2019. Before that, he was a senior researcher at Autonomous Systems Lab., ETH, Switzerland, since 2016 and obtained PhD degree in 2014 from the School of Electrical Engineering and Computer Science at QUT, Australia. His research interests lie in the development of novel approaches for Field Robotics to



highly dynamical environments. His major goal is to find new ways to deal with uncertainties and enable the design of highly interactive and adaptive systems. Prominent application examples are personal and service robots, autonomous micro-aircrafts, walking and swimming robots and driver assistant systems.

Roland Siegwart is full Professor of Autonomous Systems at ETH Zurich since July 2006 and Founding Co-Director of the Wyss Zurich. From January 2010 to December 2014 he took office as Vice President Research and Corporate Relations in the Executive Board. He received his Diploma in Mechanical Engineering in 1983 and his Doctoral Degree in 1989 from ETH Zurich. Roland Siegwart's research interests are in the design and control of systems operating in complex and



Juan Nieto is Deputy Director at the Autonomous Systems Lab, ETH Zurich, Switzerland. Before that he was a Senior Research Fellow at the Australian Centre for Field Robotics, the University of Sydney, Australia. He received his Ph.D. from the University of Sydney. His main research focus is on perception and navigation for mobile robots.