2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Hierarchical Online Learning for Adaptive Sampling of Discrete Species Distributions with an AUV

Jessica E. Todd[1], Seth McCammon[2] and Dana R. Yoerger[2]

***Abstract*— Autonomous robots are increasingly being used in the field of scientific exploration and data acquisition. In particular, the use of robotic systems for mapping and sampling of species is becoming widespread in both aerial and underwater domains, however the problem of choosing where to sample is challenging when the phenomena of interest are discrete and sparsely distributed in space or time, such as when mapping a particular benthic species. In this paper we present a hierarchical online learning framework for reasoning about species distribution in realtime, in order to inform sampling decisions. Drawing inspiration from the Species Distribution Modelling community, a hierarchical probabilistic model is developed using the Integrated Nested Laplace Approximation framework, that enables online inference about expected target hotspots using predicted substrate distributions. Model parameters are learned online to build a prediction over the discrete targets, and the model is integrated into an anytime online planner to enable adaptive path planning. The hierarchical learning approach is demonstrated on simulated synthetic environments and shown to consistently outperform baseline methods such as Gaussian Process regression and boustrophedon coverage approaches, when robot resources are constrained.**

I. Introduction

Autonomous systems are a promising tool for scalable and cost-effective exploration of the natural world [1]. In these unstructured environments, seeking and mapping the distribution of a target species of interest is critical for biodiversity monitoring and conservation [2], [3]. Species distributions are typically discrete, sparse, and exhibit high spatiotemporal variability driven by a complex interactions between abiotic and biotic processes. While these correlations are widely used by ecologists to model species distributions, they remain underutilized by conventional robotic sampling approaches, which tend to favor either open-loop lawnmower patterns [4] or adaptive sampling guided by a single scalar field [5], [6]. The key challenge for robotic systems is to make efficient use of limited sensing and energy resources to sample these heterogeneous distributions. In this paper we propose to learn and leverage ecological correlations online to enable a robot to efficiently seek out and sample species.

Species sampling presents distinct challenges compared to conventional environmental monitoring tasks. Informative path planning (IPP) has been widely explored in the context of continuous and smoothly varying fields like temperature, salinity, or chemical concentration, where gradients in the field can be directly exploited to select maximally informative sample sites [7]–[9]. In contrast, species observations are discrete and often sparsely distributed, requiring new methods for finding gradients that can be exploited by a path planner. A wealth of ecology literature demonstrates that distribution of species is largely driven by a complex interaction of biotic and abiotic factors, e.g. in the underwater context, distributions of temperature, salinity, upwelling and rainfall can predict algal blooms [10], presence of phytoplankton [11] and coral abundance [12]. In the forestry domain, abiotic factors like topography, climate and soil health combined with biotic factors like vegetation structure have been used for prediction of species richness in forests [13]. These processes also drive the broader environmental morphology, or *substrates*, in which species are found. As can be seen in Fig. 1, species are also not uniformly distributed across a given habitat, but rather exhibit spatial variation driven by both observable processes, such as the substrate distribution, and unobserved processes such as predation, social aggregation or unmodelled environmental factors.

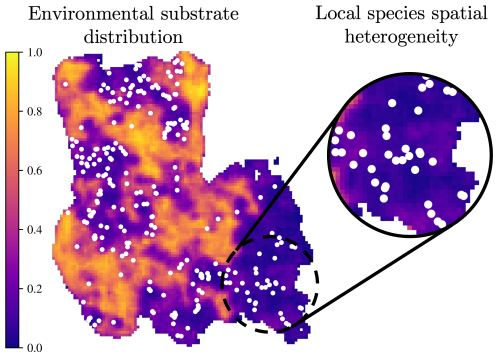

Fig. 1: Ecosystems like coral reefs can be described by substrate fields that help define reef morphology. These morphologies and other latent processes drive the presence of specific species, such as the *Porites* corals identified in this two-substrate model of Tektite Reef, U.S. Virgin Islands. The correlations between substrates (gradient field) and species (white dots) can be learned online by a robotic agent and used to guide its sampling.

We consider the problem of an autonomous underwater vehicle (AUV) tasked with locating samples of a specific coral species within a bounded region of interest. The AUV is capable of two sensing modes - a broad coverage visual sensor for detecting distributions of substrates (e.g. coral,

This work was supported in part by NSF Grants 2133029 and 2333604, and National Geographic Society Explorer Grant EC-93585T-22.

[1]Jessica Todd is with the Mechanical and Civil Engineering Department, Engineering Applied Sciences, California Institute of Technology, CA, USA jetodd@caltech.edu

[2]Seth McCammon and Dana R. Yoerger are with the Applied Ocean Physics and Engineering Department at the Woods Hole Oceanographic Institution, Woods Hole, MA, USA, {smccammon, dyoerger}@whoi.edu

rock, sand) at high altitude and a narrow field sensor for locating coral at low altitude. Unlike previous work which has considered the species sampling problem [14], we do not assume that the AUV has prior knowledge about the distribution of coral, or the ecological relationships between substrates and species. To address this gap, we introduce several key novelties: substrates are modelled as cross-correlated continuous fields, species density is assumed to be non-uniform across substrates and exhibits local variation, and the correlation between substrates and species is not known *a priori*.

To reason online about the distribution of both substrates and species, we introduce a hierarchical inference model, Gaussian Process driven Integrated Nested Laplace Approximation (GP-INLA). This hierarchical approach first models substrates using Gaussian Process (GP) regression on high-altitude observations of substrates. These predictions of substrate distributions are then used along with the low-altitude observations of species density modelled as a structured spatial point process. Integrated Nested Laplace Approximations (INLA) is used to infer the relationship between the observed species point process and the predicted substrate covariates, allowing the robot to produce a belief map of species density in real-time during a species sampling mission. Finally, this hierarchical model is integrated into an adaptive planning framework that balances online learning of the GP-INLA model with locating target species. The main contributions of this work are (1) A hierarchical generative model for discrete species distributions over multiple substrate types and latent fields, (2) an informative path planning algorithm that integrates our hierarchical model into an adaptive planning framework and (3) experiments using synthetic datasets of coral distributions demonstrating both a significant reduction in RMSE over the duration of the mission, and up to a 70% performance increase in seeking out unique coral targets compared to boustrophedon coverage and a GP-based adaptive planning framework that does not explicitly reason about the species-substrate relationship.

II. Related Work and Background

Recently, the use of environmental covariates, or substrates, has been explored to model hard-to-observe fields. Proxy measurements of chlorophyll concentration, salinity and temperature have been used to maximise plankton sampling across multiple missions [11]. Bayesian networks have also been leveraged to map sensor data and scientific prior knowledge of geologic variables of interest as part of a Monte Carlo Tree Search (MCTS) framework [15], [16]. Spatial correlation is loosely enforced through a simple Gaussian model that allows limited prediction of unseen regions. In contrast to these approaches, we use GPs fitted to our substrates to enable non-myopic planning over long horizons, and our hierarchical approach is able to reason online about the distribution of our target of interest. However, species modelling techniques take these methods a step further, leveraging point process models to capture the complex spatial interactions betwen a species distribution and its biotic and abiotic drivers [17]–[19].

A Poisson point process [20] describes a stochastic point process where for a measurable region $A \subseteq \mathbb{R}^d$, the number of points $N(A)$ are Poisson distributed with mean $\int_A \lambda(x)dx$ and conditional on $N(A)$ are i.i.d. with a density proportional to the intensity function $\lambda(x)$. A Cox point process is a generalized form of the Poisson point process where the intensity function is itself a random function $\Lambda(x)$, and points are conditional on a random intensity field sampled from $\Lambda(x)$. The log Gaussian Cox process (LGCP) [21] is a special case of the Cox process that models the random intensity field as a $\Lambda(x) = exp(Z(x))$ where $Z(x)$ is a Gaussian process that allows for the introduction of spatial correlation structure via its covariance function.

Species modelling takes the LGCP one step further by relaxing the local conditional independence assumption of the Poisson point process, instead modelling the intensity as a loglinear combination of contributing and observable environmental spatial data, $\mathbf{z}(x)$ or *covariates*, and a realization of a random field $\zeta(x)$

$$\log \lambda(x) = \alpha_\lambda + \mathbf{z}^T(x)\beta_\lambda + \zeta(x). \tag{1}$$

An example of the LGCP is shown in Fig. 2. Environmental covariates capture measured fields such as temperature, rainfall etc. The assumption is that the introduction of the random effects term $\zeta(x)$ captures the remaining unmeasured or latent processes associated with the distribution of a species [22]. In practice it is not feasible to observe and measure all fields that might influence the species distribution, and indeed there are other inter-species effects such as predation or social aggregation which can drive localised heterogeneity in the species distribution. In robotic search, LGCPs have been used to model the distribution of survivors in a search-and-rescue scenario [23], where Integrated Nested Laplace Approximations (INLA) are used to learn the relationship between survivor density and risk factors in real time as people are detected. Like prior species search methods [14], a key assumption of Andersson's approach is access to *a priori* environmental data across the entire region of interest, such as building distribution, which provides strong priors to each of the models, an assumption which does not hold in many real-world environments.

III. Problem Formulation

Our problem formulation follows the Sparse Search-and-Sample Problem definition proposed by Todd et. al [14], which formulates sparse-search-and-sample using a partially-observable Markov decision process (POMDP) [24]. For brevity we do not reproduce the full POMDP definition, but rather describe the modifications to it to account for the relaxation of assumptions on *a priori* knowledge of substrate-species covariation. In this paper, we make two modifications to the POMDP in [14]. First we modify the state space, $\mathcal{S}$, expanding the world state to include the substrate state, $v \in V = \mathbf{W}^{m \times n}$, which is a k dimensional vector of each cell comprised of a mixture of k substrate distributions

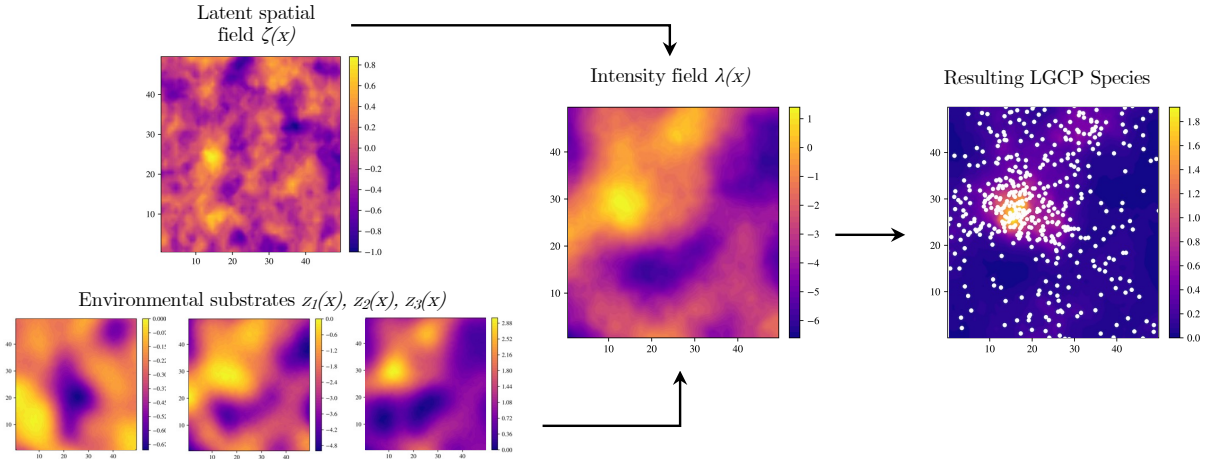

Fig. 2: LGCP Generative Species Model. Known substrate covariates $\mathbf{z}(x)$ are combined with the latent spatial field, $\zeta(x)$ (which accounts for the remaining effects of all covariates not explicitly modelled in $\mathbf{z}$), as a weighted sum as described in (1) to form the LGCP intensity field, $\lambda(x)$. Samples drawn from this field form the species distribution.

$\mathbf{W} = \{f_1, ...f_k\}$ $s.t.$ $\sum f_i = 1$. Secondly, we include the ability to observe these distributions of substrates, v, in the observation space $\mathcal{O}$ when the AUV is at high altitude as well as the ability to count the number of target species, u, observed at low altitude.

We formulate the optimal exploration for species sampling as an IPP, where we seek a trajectory $\phi = [x_1, x_2, ...x_B]$ in the space of all possible trajectories Φ that maximises a reward function, R, subject to a resource budget constraint, B, on the cost of the path, C, such as energy or time.

$$\phi^* = \arg\max_{\phi \in \Phi} I(\phi) \text{ s.t. } C(\phi) \leq B \tag{2}$$

Since the proposed approach necessitates online learning of a species-substrate model, the information-theoretic measure can be reframed as an acquisition function I that combines both the decision-theoretic maximisation of reward for observing a species,

$$R(\phi) = \begin{cases} \sum_{\phi} u(x) \text{ if } x \in x_{low} \\ 0 \qquad\qquad \text{otherwise} \end{cases} \tag{3}$$

and an information-theoretic reduction of uncertainty within the species distribution model.

IV. Approach

GP-INLA combines two modelling approaches into a unified hierarchical model for discrete sample distribution. At the top level, we use a GP mixture to model the distribution of different substrate types within our region of interest. We assume that the distribution of these substrates correspond to environmental morphologies that are observed at high altitude. We model the species as a LGCP conditioned on both the substrate distribution, and a latent gaussian field to capture local spatial variation in the species distribution, and use INLA to determine the expected density of the target species over the region. Finally, we integrate this hierarchical model in an adaptive planning algorithm designed to balance optimizing model parameters with maximizing the number of observations of the target species using an anytime Monte Carlo Tree Search (MCTS) planner. This pipeline is shown in Fig. 3.

A. Gaussian Process regression

The LGCP model requires environmental substrate data over the entire region of interest. Even if covariate data are not available at specific locations where species are observed, it must be sampled densely enough across the world that interpolation can be made between species observation locations to enable predictions of the intensity field. Since the world is partially observable, substrates are only observable within the footprint of the sensor. To counter this limitation, we propose the use of GP regression models fit to each substrate, to provide a prediction over the substrate distribution.

A GP is a distribution over functions that can be fully specified by a mean function and covariance function $f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), K(\mathbf{x}, \mathbf{x}'))$ [25]. In the GP-INLA framework, observations of the substrates are taken by the robot at high altitude. These observations are a vector of substrate distributions in a given grid cell, at location x visited by the robot, normalised to one. To ensure the GP is zero-mean, observations for each substrate are shifted to be the deviation from the mean value of the substrate as observed so far. A Radial Basis Function kernel (RBF) was selected as the covariance function. Substrate observations and observation locations are passed to the GP module as training points, and the hyperparameters (ℓ, σ) of the kernel function are periodically trained. The mean of these distributions are passed to the INLA model as the current estimate of the substrate distribution, and species inference conducted based on these values. GP regression has the advantage of fully characterising the uncertainty in the substrate data through the variance term σ, a parameter which can be exploited in the designing of acquisition functions to motivate exploration in unknown regions of the world. Our GP framework was implemented using the `GPy` library [26].

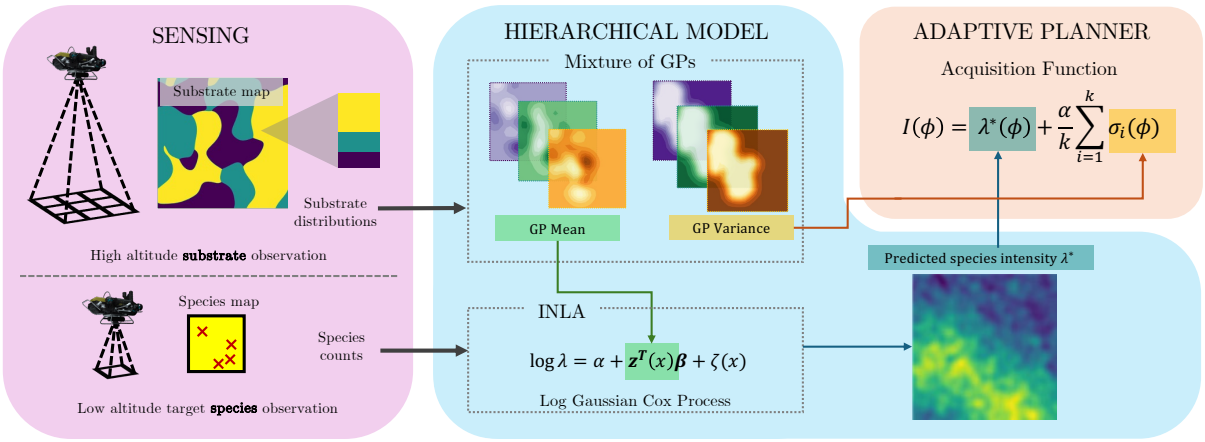

Fig. 3: Diagram of the GP-INLA framework. The robot can sense at either high altitude, detecting substrate distributions v or at low altitude detecting species observations u. Substrate distribution $v = [f_1, f_2, ...f_k]$ is passed to the top level of the hierarchical model, and a mixture of k GPs used to model the substrate distributions over the region of interest. The GP means $[\mu_1, \mu_2, ...\mu_k]$ are passed to the lower level of the GP-INLA model as estimates over the predicted environmental substrate distributions used in the LGCP model, and the uncertainty $[\sigma_1, \sigma_2, ...\sigma_k]$ are passed directly to the acquisition function. Using the low altitude species observations u and substrate distributions, INLA performs Bayesian inference over variables of the LGCP model $(\alpha, \boldsymbol{\beta}, \zeta)$ to produce a posterior of the predicted species intensity λ. Both λ and σ_i are used in the acquisition function to reason over the most informative path ϕ.

B. Online Learning using INLA

Inference over our species distribution is performed using Integrated Nested Laplace Approximation (INLA [27]). If we model the species as a realisation of an LGCP, then the number of observed target species u_i at a given location x_i is given by

$$u(x) \sim \text{ Poisson}(\lambda(x)),\ \lambda(x) = \exp(\eta(x)) \tag{4}$$

where η represents the log link function. Environmental substrate data $\mathbf{z}$ is taken as the mean μ_i where $i = 1, ..., k$ from the substrate GPs as the current estimate of the substrate fields.

Exact Bayesian inference such as Markov Chain Monte Carlo (MCMC) over the latent field $\zeta(x)$ and coefficients α and β is computationally expensive. INLA is an approximate Bayesian inference method that is fast and memory efficient due to its use of nested Laplace approximations to compute approximate posteriors for the coefficients and latent field, and leveraging sparse precision matrices by representing the latent spatial field as a GMRF on a mesh. INLA is a standard method for spatial analysis in statistical and ecology literature but is rarely applied to Robotics.

If we consider the LGCP model described by (1), let $\boldsymbol{\psi}$ be a vector of Gaussian variables containing the coefficients $\alpha, \boldsymbol{\beta}$ and the discretized random field ζ, with prior

$$\boldsymbol{\psi} = (\alpha, \boldsymbol{\beta}, \zeta) \sim \mathcal{N}(\mathbf{0}, Q(\boldsymbol{\theta})^{-1}) \tag{5}$$

where $Q(\boldsymbol{\theta})^{-1}$ is the sparse precision matrix and $\boldsymbol{\theta} \sim \pi(\boldsymbol{\theta})$ is a vector of hyperparameters describing the Matern kernel for ζ. INLA approximates the posterior

$$\pi(\boldsymbol{\psi}, \boldsymbol{\theta}|\mathbf{u}) \propto \pi(\mathbf{u}|\boldsymbol{\theta}, \boldsymbol{\psi})\pi(\boldsymbol{\psi}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}) \tag{6}$$

The hyperparameter posterior is approximated using Laplace approximations

$$\pi(\boldsymbol{\theta}|\mathbf{u}) \approx \tilde{\pi}(\boldsymbol{\theta}|\mathbf{y}) = \frac{\pi(\mathbf{x}, \theta, \mathbf{y})}{\tilde{\pi}_G(\mathbf{x}|\boldsymbol{\theta}, \mathbf{y})}\Big|_{\mathbf{x}=\mathbf{x}^*(\boldsymbol{\theta})} \tag{7}$$

where $\tilde{\pi}_G$ is a Gaussian approximation to the full conditional posterior of $\boldsymbol{\psi}$ and $\boldsymbol{\psi}^*$ is the mode of $\boldsymbol{\psi}$ for a given set of hyperparameters $\boldsymbol{\theta}$. The marginal posteriors of each variable ψ_i

$$\pi(\psi_i|\mathbf{u}) = \int \pi(\psi_i|\mathbf{u}, \boldsymbol{\theta})\pi(\boldsymbol{\theta}|\mathbf{u})d\boldsymbol{\theta} \tag{8}$$

is obtained by marginalising a Laplace approximation of $\pi(\psi_i|\mathbf{u}, \boldsymbol{\theta})$ and integrated numerically over a low-dimensional grid of $\boldsymbol{\theta}$. Thus a simple description of the INLA pipeline is (1) the posterior marginal of $\boldsymbol{\theta}$ is computed using the Laplace approximation, (2) $\pi(\psi_i|\mathbf{u}, \boldsymbol{\theta})$ is computed using the Laplace approximation for select values of $\boldsymbol{\theta}$, and (3) the previous two approximations are combined using numerical integration. For a detailed explanation of the INLA approach we direct the interested reader to [27].

If we take x^* as unobserved locations in the region of interest, then using the marginals calculated in (8) can be used to get the posterior mean of the species intensity $\mathbb{E}[\lambda(x^*)|\mathbf{u}]$, by integrating over

$$\mathbb{E}[\lambda(x^*)|\mathbf{u}] = \int \exp(\eta)\pi(\eta|\mathbf{u})d\eta \tag{9}$$

where

$$\pi(\eta(x^*)|\mathbf{u}) = \int \pi(\eta(x^*)|\mathbf{u}, \boldsymbol{\theta})\pi(\boldsymbol{\theta}|\mathbf{u})d\boldsymbol{\theta} \tag{10}$$

and $\eta(x) = \log\lambda(x) = \alpha + \mathbf{z}(x)^T\boldsymbol{\beta} + \zeta(x)$.

Using the INLA method we obtain a predicted species intensity $\lambda^*(x)$. Our INLA framework was implemented

using the R-INLA library in R [28]. We use the regular lattice meshing method in INLA, and employ the Empirical Bayes (EB) estimator during the inference process to speed up estimation of the model parameters.

C. Anytime Adaptive Path Planning

The formulation of our formal Search-and-Sample problem means there is no direct reward obtained from observing substrates, (see (3)) yet the robot must be motivated to explore the substrate fields in order to improve the model of species intensity. As is described in the IPP formulation we employ an informative acquisition function $I(\cdot)$ to motivate the robot to both explore new regions of the map and exploit areas of known high species density. We based our acquisition function on the classic UCB reward function to balance exploration and exploitation. However our case differs from the classic UCB case as the exploitative utility and explorative utility are being drawn from two different parts of the model. Due to the stochasticity in the species observation data, calculated variance in the predicted species distribution produced by INLA captures both aleatoric uncertainty (variability in the data) and epistemic uncertainty (uncertainty in the model). Consequently we see high variance in the predicted species distribution, even in areas that have already been observed by the robot, as the inference model attempts to fit a smooth predictive intensity to stochastic data. Parsing out the epistemic uncertainty and aleatoric from this model during the planning phase of a mission, in an effort to make decisions to reduce epistemic model uncertainty, would require selecting a new observation point, sampling the predicted species density at that location, adding this back into the current set of observations, and rerunning INLA to assess how the variance changes with the model, for each proposed observation along a potential path. With a computational complexity of $O(n^{3/2})$ this process is infeasible in an online setting. Consequently, the variance, or uncertainty, output by the INLA module is not utilised in this acquisition function.

Instead we opt to use heuristics that allow for approximation of the uncertainty in the model. The exploit term comes directly from the predicted species intensity, produced by the full hierarchical model, λ^*. The exploration term comes from the variance, or uncertainty, in the substrate model, as this characterised areas of the region of interest with low confidence in the substrate distribution, and thus the accuracy of the INLA results at these locations will also have high uncertainty. Since we have k substrates, we have $\sigma = [\sigma_1, \sigma_2, ...\sigma_k]$. Thus the overall acquisition function is given by

$$I(\phi) = \lambda^*(\phi) + \frac{\alpha}{k} \sum_{i=1}^{k} \sigma_i(\phi). \tag{11}$$

A fixed-horizon MCTS planner [29] was used for the sequential decision making process, with the above acquisition function evaluated for each proposed rollout and backpropagated through the decision tree. The world state (both substrate and species) is estimated using GP-INLA, and partial actions forward simulated to a horizon h over 5 possible actions (4-connected grid plus an altitude change). Actions that move the robot into a previously unexplored grid cell are randomly selected from first, and repeat cell visits are considered only if there are no other viable actions. For a complete description of MCTS we refer the reader to [29].

V. Experimental Results

We assessed the empirical performance of the hierarchical adaptive planner in simulations using synthetically generated environments. We first assessed the suitability of the hierarchical approach in approximating the ground truth distribution of the species, without the integration of the adaptive planner. We then assessed the performance of the adaptive planning framework in seeking out and sampling species targets, compared to baseline methods.

A. Data generation

Synthetic worlds were generated using a generative model inspired by the LGCP. Since the goal is to make inference about a world that exploits both natural cross-substrate correlation and spatial correlation of substrates, the synthetically-generated data must exhibit these relationships. Substrates are simulated as spatially distributed Gaussian random fields. Each of the k substrates is drawn from a zero-mean GP prior based on the squared exponential covariance function, with hyperparameters (ℓ, σ_n^2). For each world, these lengthscale and variance parameters were kept constant across substrates. To ensure spatial cross-correlation between substrates, the final 'ground-truth' substrate fields are taken as joint distributions resulting from different combinations of these initial fields. Random Gaussian noise $\epsilon \sim \mathcal{N}(0, 0.05^2)$ is then added to each field. A Gaussian Markov random field (GMRF) is used to simulate unmeasured covariates and latent behaviours driving localised heterogeneity, using a Matern covariance function with $\nu = 1$. The final species point process is a realisation of this simulated intensity field given in (1), where $\mathbf{z}$ is a vector of joint substrate distribution fields and ζ is the GMRF. Fixed coefficients $(\alpha, \beta_{\lambda_1}, \beta_1, ...\beta_k)$ were randomly selected to be between $(-4, 6)$, and then tuned lightly to ensure a reasonable number of species targets in each world. β coefficients were selected such that each world had a mix of substrates that were positively and negatively correlated with species presence.

B. Comparison Against Benchmarks

The proposed hierarchical model and planner were evaluated against two baseline algorithms. The first is a traditional coverage planner performing a boustrophedon (lawnmower) pattern over the world. This non-adaptive approach operates at low altitude, at a velocity of 1 m/s taking a total of 2500 seconds to map the entire 50 $\times$ 50 m world. The coverage planner was simulated running both vertical and horizontal tracklines from each of the four corners of the world, for a total of 8 missions per world. The second baseline is a GP regression model-based planner operating directly on the species observations. The GP model takes

the number of species in a given grid cell as the direct training input to the model, and then makes predictions over the remaining unvisited grid cells. Like GP-INLA, the GP model is integrated into a receding horizon MCTS planner with identical parameters as GP-INLA, with the world state estimate provided by the predicted mean from the GP. The acquisition function used for the GP approach was

$$I(\phi) = \mu(\phi) + \alpha \times \sigma(\phi) \tag{12}$$

where μ and σ are the mean and variance of the predicted species distribution field from the GP model. Through empirical evaluation, $\alpha = 0.15$ was found to produce the best performance.

C. Evaluation of the Hierarchical Model

We first evaluated the performance of the hierarchical model in correctly approximating the species distribution. We compare against the baseline approach of fitting a nonhierarchical GP regression model directly to the species observations. As we want to assess prediction of the underlying intensity process from which the species is a realisation, without factoring in online adaption, the robot moves in a simple boustrophedon coverage trajectory and employs either the GP or GP-INLA method to generate a current belief map of the world. The robot moves from grid cell to grid cell, observing the number of species in each cell. In the case of the GP-INLA model, which takes both substrate and species observations, a high-altitude substrate observation is incorporated every 10 steps, with a transition time of 5 seconds. In both the GP and GP-INLA case, the model is retrained every 10 steps to produce an updated estimate of species distribution. For each world, the robot performed 8 missions for each of 3 trackline densities (1, 2 and 5 m spacing) with sensor footprints held constant, to assess how the model performs with only partial data. Since INLA performs poorly with few spatially-concentrated datapoints, both GP and GP-INLA were warm-started using observations collected along a bounding-box trajectory around the region, with periodic altitude sampling for GP-INLA.

Fig. 4 shows the results for a three-substrate world, with predictive ability assessed using the total Root Mean Squared Error (RMSE) between the ground truth intensity function λ and the predicted intensity function λ^* across all grid cells. In all three coverage densities, we see an improvement by GP-INLA over the standard GP model, with a peak decrease in RMSE of $41.9\%, 41.6\%$ and 39.1% for 1m, 2m and 5m respectively. The GP-INLA model is initially very noisy, seen as a large spike in the RMSE value. This is a consequence of the model fitting poorly with extremely limited data, however it is able to recover quickly and the sharp slope in the GP-INLA plot shows that the model is quickly learning an estimate of λ. GP-INLA outperforms the GP model in all three cases until right at the end of the mission, and significantly outperforms GP in the sparsest trackline. The sparse observations together with the stochastic data, aren't well captured by the GP alone, as the GP is attempting to fit a smoothly varying field to an inherently stochastic set of patchy observations. The lengthscale hyperparameter is being over- or underfit to the stochastic species observations, resulting in poor estimation of species distribution. In contrast, the GP-INLA model factors this stochasticity into the model, as well as the smoother substrate information.

D. Evaluation of an Adaptive Online Learning Planner

In this experiments, we evaluated the full adaptive planning pipeline for GP-INLA, in adaptively trying to maximise unique observations of target species. We compare against the single low-altitude adaptive GP approach, and a low-altitude non-adaptive boustrophedon coverage planner. The low altitude GP model was retrained every 10 seconds of mission time. For GP-INLA, the INLA model was retrained every 10 seconds of mission time and the substrate GP retrained every time a new high altitude observation of substrates was taken. We generated 5 synthetic worlds, one of which is shown in Fig. 4, and the remaining 4 in Fig. 5a, with 8 missions per world for a total of 40 simulations. In each case the robot starts in the corner of the world and takes an initial boundary trajectory for gathering the warm-start observations. The resulting reward performance (unique target species observed) is shown in Fig. 5b, as a function of mission time. The missions were run for 2500s, including the boundary path, the time it takes for a coverage planner to complete full coverage of the world.

Across all worlds, GP-INLA demonstrates statistically significant performance improvements over the first 2000 s of mission time, as evidenced by the lack of overlap in the standard error bars, demonstrating both a faster rate of reward acquisition in the first half of the mission, and a greater overall reward gathered until 2000s. The GP approach performs extremely poorly in all but one world, underperforming even the boustrophedon planner for the majority of the mission. Again, the GP model fails to properly capture the smoothly varying distribution of species from sparse species observation, making it ill-suited to the discrete search and sample problem. World 1 exhibits a clear hotspot with more of a gradient of observations, particularly along the edge of the world, which may explain why the performance of the GP algorithm was slightly better at capturing reward in this instance. World 3 shows the smallest performance difference between GP-INLA and the boustrophedon pattern. This world is characterised by a far more homogeneous spread of reward than the other worlds, and so the coverage planner will do a good job of gathering this reward over the course of the mission. GP-INLA is well-suited to instances where there is evident hotspots or clustering that can be exploited, as is more often the case with real coral reef data. GP-INLA exhibits similar performance in all cases, moving towards predicted hotspots early on in the mission, while continuing to gather data about the world and thus the species distribution, resulting in a sharp rate of reward initially. As these hotspots are exploited, the planner drives the robot towards unknown areas within the map to continue gathering reward. Given sufficiently long mission time, the coverage algorithm will eventually outperform the adaptive

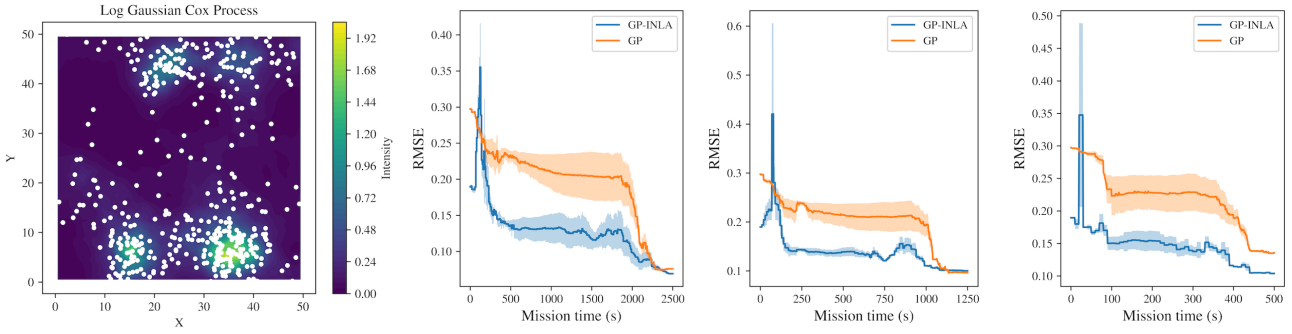

(a) True world state (b) 1m tracklines (c) 2m tracklines (d) 5m tracklines

Fig. 4: Root Mean Squared Error (RMSE) of a $k = 3$ substrate world. *(a)* shows the true $\log \lambda$ intensity field and generated species point process. *(b)-(d)* show the RMSE between the predicted spatial distribution from the inference model, and the true intensity field from which the point process was generated for 1m, 2m and 5m spaced tracklines respectively. The shaded area is standard error

algorithm over the same fixed mission time, as the coverage algorithm visits every cell sequentially, and is not permitted to backtrack, whereas the GP-INLA approach will initially ignore areas of predicted sparsity, and thus can potentially need to retread visited cells towards the end of the mission to gather all possible reward. GP-INLA is therefore the better choice when operating in a restricted mission budget (i.e., budgeted mission time is much less than time required for full coverage).

VI. Conclusion and Future Work

In this paper we developed a hierarchical approach to mapping discrete and inhomogenously-distributed species in an unknown environment. Our approach, GP-INLA, uses simultaneous inference with Gaussian Process regression to map substrates and integrated nested Laplace approximations to quantify the relationship between substrates and the target species. Comparing the GP-INLA approach to other baseline methods for mapping sparse targets, GP-INLA is able to more quickly produce an accurate estimate of the true species distribution function. In simulated experiments we showed that using the additional context provided by explicitly reasoning over substrate types, GP-INLA was able to achieve up to a 70% improvement over a standard GP. Compared to exhaustive lawnmower coverage, GP-INLA achieves asymptotic performance in observing all of the target species individuals, while observing more individuals in the first $70 - 80\%$ of each mission.

A clear avenue for future work is to adapt the proposed approach for real world deployment using a low cost AUV. Species detections can be achieved with a realtime object detection algorithm such as a YOLO detector [30] trained on a specific coral species, with substrate observations coming from a semantic classifier such as Realtime Online Spatiotemporal Topic modelling [31], as demonstrated in [14]. While results for this paper were generated on a compute server to accelerate experimentation, the same implementation was demonstrated to run in real time on a System76 Oryx Pro laptop with Intel i7 CPU and 64 GB RAM.

Additionally, incorporating the additional uncertainty of the model into the acquisition function, by examining the impact of epistemic uncertainty in the INLA component of this model, could help guide the robot to highly valuable observations that would improve the estimates of the LGCP model parameters. While this work focused on the underwater domain, which introduces certain mission limitations such as robot speed at different altitude and sensor degradation with distance, this work can be more broadly applied to other scientific domains.

References

[1] M. Dunbabin and L. Marques, "Robots for Environmental Monitoring: Significant Advancements and Applications," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 24–39, 2012.

[2] K. Shah, G. Ballard, A. Schmidt, and M. Schwager, "Multidrone aerial surveys of penguin colonies in Antarctica," *Science Robotics*, vol. 5, no. 47, p. eabc3000, 2020.

[3] W. Reckling, H. Mitasova, K. Wegmann, G. Kauffman, and R. Reid, "Efficient Drone-Based Rare Plant Monitoring Using a Species Distribution Model and AI-Based Object Detection," *Drones*, vol. 5, no. 4, p. 110, 2021.

[4] S. B. Williams, O. Pizarro, M. Jakuba, and N. Barrett, "Auv benthic habitat mapping in south eastern tasmania," in *Field and Service Robotics*, A. Howard, K. Iagnemma, and A. Kelly, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 275–284.

[5] S. McCammon, G. Marcon dos Santos, M. Frantz, T. P. Welch, G. Best, R. K. Shearman, J. D. Nash, J. A. Barth, J. A. Adams, and G. A. Hollinger, "Ocean front detection and tracking using a team of heterogeneous marine vehicles," *Journal of Field Robotics*, vol. 38, no. 6, pp. 854–881, 2021.

[6] K.-C. Ma, L. Liu, H. K. Heidarsson, and G. S. Sukhatme, "Data-driven learning and planning for environmental sampling," *Journal of Field Robotics*, vol. 35, no. 5, pp. 643–661, 2018.

[7] G. Hitz, E. Galceran, M.- Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017.

[8] M. Popović, T. Vidal-Calleja, G. Hitz, J. J. Chung, I. Sa, R. Siegwart, and J. Nieto, "An informative path planning framework for UAV-based terrain monitoring," *Autonomous Robots*, vol. 44, no. 6, pp. 889–911, 2020.

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

[10] C. Anderson, M. Sapiano, M. Prasad, W. Long, P. Tango, C. Brown, and R. Murtugudde, "Predicting potentially toxigenic Pseudo-nitzschia blooms in the Chesapeake Bay," *Journal of Marine Systems*, vol. 83, no. 3-4, p. 14, 2010.

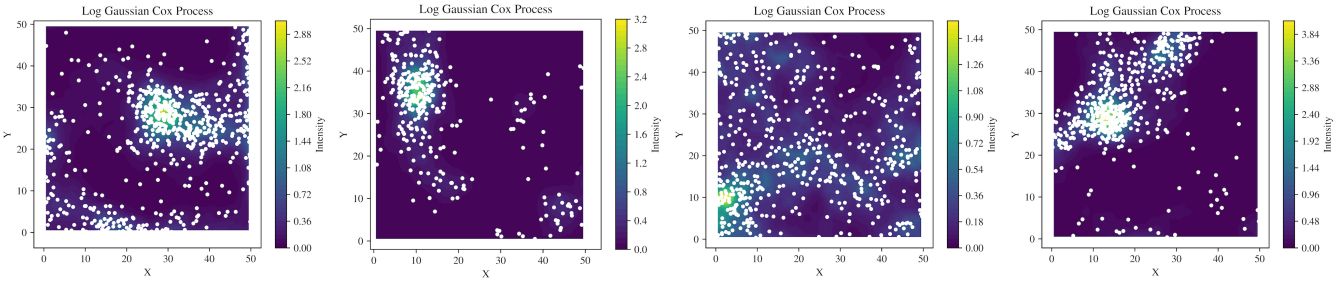

(a) Species intensity field for synthetic worlds 1-4.

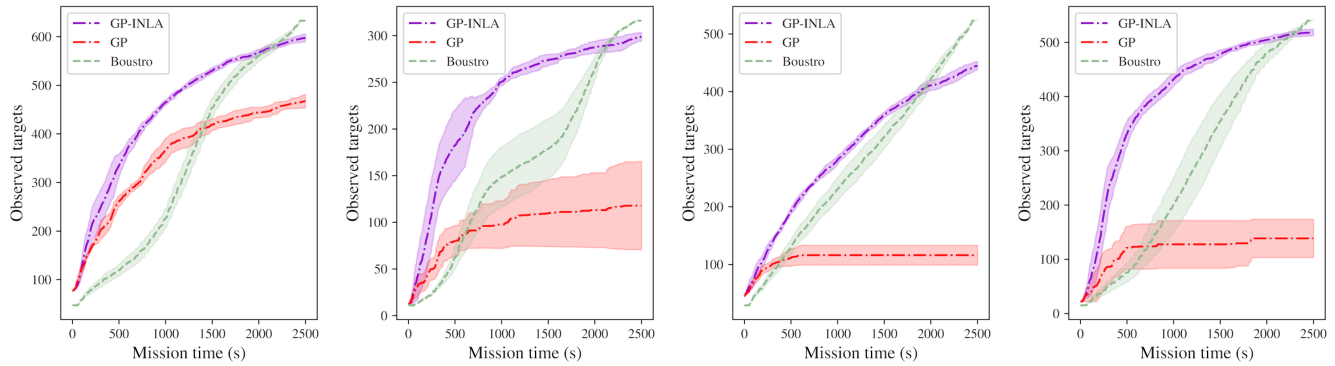

(b) Mean $\pm\sigma_N$ number of unique targets observed for $N = 8$ missions.

Fig. 5: Comparison of GP-INLA to GP and boustrophedon. *Top row:* Synthetic worlds 1-4 showing the species intensity and species presence. All worlds have $k = 3$ substrates with randomly selected parameters, and species presence are shown in white. *Bottom row:* Performance in terms of mean targets observed for $N = 8$ missions on worlds 1-4. Note that worlds and corresponding performance for that world are aligned vertically. Results for world 5 is omitted due to space constraints.

[11] J. Das, F. Py, J. Harvey, J. Ryan, A. Gellene, R. Graham, D. Caron, K. Rajan, and G. Sukhatme, "Data-driven robotic sampling for marine ecosystem monitoring," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1435–1452, 2015.

[12] S. Georgian, K. Kramer, M. Saunders, W. Shedd, H. Roberts, C. Lewis, C. Fisher, and E. Cordes, "Habitat suitability modelling to predict the spatial distribution of cold-water coral communities affected by the Deepwater Horizon oil spill," *Journal of Biogeography*, vol. 47, pp. 1–12, 2020.

[13] F. Zellweger, A. Baltensweiler, C. Ginzler, T. Roth, V. Braunisch, H. Bugmann, and K. Bollmann, "Environmental predictors of species richness in forest landscapes: abiotic factors versus vegetation structure," *Journal of Biogeography*, vol. 43, no. 6, pp. 1080–1090, 2016.

[14] J. E. Todd, "Adaptive Robotic Search and Sampling of Sparse Natural Phenomena," Thesis, Massachusetts Institute of Technology, May 2024.

[15] A. Arora, R. Fitch, and S. Sukkarieh, "An approach to autonomous science by modeling geological knowledge in a Bayesian framework," in *2017 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2017, pp. 3803–3810.

[16] A. Arora, P. M. Furlong, R. Fitch, S. Sukkarieh, and T. Fong, "Multi-modal active perception for information gathering in science missions," *Autonomous Robots*, vol. 43, no. 7, pp. 1827–1853, 2019.

[17] W. J. Wright, K. M. Irvine, T. J. Rodhouse, and A. R. Litt, "Spatial Gaussian processes improve multi-species occupancy models when range boundaries are uncertain and nonoverlapping," *Ecology and Evolution*, vol. 11, no. 13, pp. 8516–8527, 2021.

[18] N. Golding and B. V. Purse, "Fast and flexible Bayesian species distribution modelling using Gaussian processes," *Methods in Ecology and Evolution*, vol. 7, no. 5, pp. 598–608, 2016.

[19] I. W. Renner, J. Elith, A. Baddeley, W. Fithian, T. Hastie, S. J. Phillips, G. Popovic, and D. I. Warton, "Point process models for presence-only analysis," *Methods in Ecology and Evolution*, vol. 6, no. 4, pp. 366–379, 2015.

[20] J. Møller and R. P. Waagepetersen, "Modern Statistics for Spatial Point Processes," *Scandinavian Journal of Statistics*, vol. 34, no. 4, pp. 643–684, 2007.

[21] J. Møller, A. R. Syversveen, and R. P. Waagepetersen, "Log Gaussian Cox Processes," *Scandinavian Journal of Statistics*, vol. 25, no. 3, pp. 451–482, 1998.

[22] P. J. Diggle, P. Moraga, B. Rowlingson, and B. M. Taylor, "Spatial and Spatio-Temporal Log-Gaussian Cox Processes: Extending the Geostatistical Paradigm," *Statistical Science*, vol. 28, no. 4, pp. 542–563, 2013.

[23] O. Andersson, P. Sidén, J. Dahlin, P. Doherty, and M. Villani, "Real-Time Robotic Search using Structural Spatial Point Processes," in *35th Uncertainty in Artificial Intelligence Conference*, 2020, pp. 995–1005.

[24] S. Russell and P. Norvig, *Artificial intelligence: a modern approach*. Upper Saddle River, NJ, USA: Prentice-Hall, 2002.

[25] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, F. Bach, Ed. Cambridge, MA, USA: MIT Press, 2005.

[26] GPy, "GPy: A gaussian process framework in python," http://github.com/SheffieldML/GPy, 2012.

[27] H. Rue, S. Martino, and N. Chopin, "Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 71, no. 2, pp. 319–392, 2009.

[28] R Core Team, *R: A Language and Environment for Statistical Computing*, https://www.r-inla.org, R Foundation for Statistical Computing, 2025.

[29] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, "A Survey of Monte Carlo Tree Search Methods," *IEEE Transactions on Computational Intelligence & AI in Games*, vol. 4, no. 1, pp. 1–43, 2012.

[30] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 779–788.

[31] Y. Girdhar, P. Giguère, and G. Dudek, "Autonomous adaptive exploration using realtime online spatiotemporal topic modeling," *The International Journal of Robotics Research*, vol. 33, no. 4, pp. 645–657, 2014.