

Informative Path Planning to Estimate Quantiles for Environmental Analysis

Isabel M. Rayas Fernández^{†,1}, Christopher E. Denniston^{†,1}, David A. Caron², Gaurav S. Sukhatme^{‡,1}

Abstract—Scientists interested in studying natural phenomena often take physical specimens from locations in the environment for later analysis. These analysis locations are typically specified by expert heuristics. Instead, we propose to choose locations for scientific analysis by using a robot to perform an informative path planning survey. The survey results in a list of locations that correspond to the quantile values of the phenomenon of interest. We develop a robot planner using novel objective functions to improve the estimates of the quantile values over time and an approach to find locations which correspond to the quantile values. We test our approach in four different environments using previously collected aquatic data and validate it in a field trial. Our proposed approach to estimate quantiles has a 10.2% mean reduction in median error when compared to a baseline approach which attempts to maximize spatial coverage. Additionally, when localizing these values in the environment, we see a 15.7% mean reduction in median error when using cross-entropy with our loss function compared to a baseline.

Index Terms—Field Robots, Planning under Uncertainty, Marine Robotics, Environment Monitoring and Management

I. INTRODUCTION

IN order to understand biological phenomena, scientists take a small number of physical specimens for later analyses to characterize the biological community and the contextual environmental conditions at the site. A marine biologist may capture water in a container or filter media at a location for later analysis, and an agricultural scientist may capture a small portion of a plant. Scientists later analyze these captured portions of the environment in a detailed laboratory setting. It is often expensive to collect these physical specimens as the scientist must go to the location and use a physical reagent, such as a container or filter media, which are limited in the field. Traditionally, expert heuristics underpin the selection

Manuscript received: February 24, 2022; Revised May 25, 2022; Accepted June 29, 2022.

This paper was recommended for publication by Editor Hyungpil Moon upon evaluation of the Associate Editor and Reviewers' comments.

This work was supported in part by the Southern California Coastal Water Research Project Authority under prime funding from the California State Water Resources Control Board on agreement number 19-003-150 and in part by USDA/NIFA award 2017-67007-26154.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1842487. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

[†] Equal contribution. All authors are with the University of Southern California. ¹ Robotic Embedded Systems Lab, ² Caron Protistan Aquatic Microbial Ecology Lab {rayas, cdennist, dcaron, gaurav}@usc.edu

[‡] G.S. Sukhatme holds concurrent appointments as a Professor at USC and as an Amazon Scholar. This paper describes work performed at USC and is not associated with Amazon.

Digital Object Identifier (DOI): see top of this page.

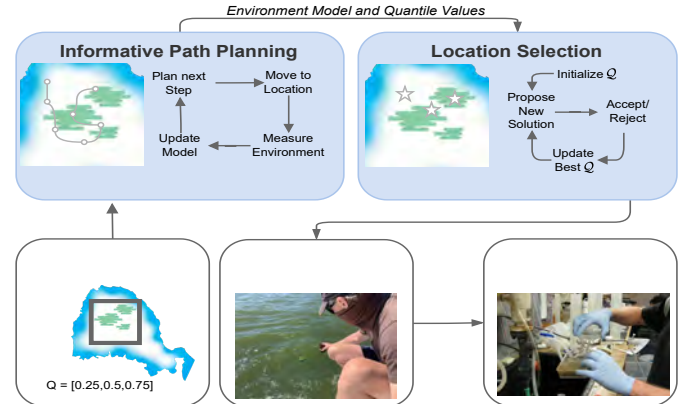


Fig. 1: **Full system.** First, the parameters for the robotic survey are chosen, such as the area bounds and the quantiles for specimen collection. The robot performs informative path planning using our proposed objective functions, creating an environment model and an estimate of the quantile values. The quantile locations are then selected by minimizing our proposed loss function. These quantile locations correspond to physical locations which, when measured, have the estimated quantile values. After the quantile locations are chosen, humans go to them to collect field specimens which they later analyze in a laboratory setting. This work focuses on the steps shown in the blue shaded boxes.

of locations for scientific analysis. These expert heuristics generally attempt to spread out the analysis locations in the phenomena of interest in order to take specimens with differing concentrations for characterization of heterogeneity [1]. The heuristics may make use of robotic surveys, but currently most robotic surveys use either pre-programmed paths [1], or autonomy which seeks maxima [2] or to cover the area spatially [3], [4], and do not directly plan for the specimen collection tasks. In contrast, we propose performing an adaptive robotic survey to find locations of interest for scientific analysis. To specify these locations, we propose calculating the quantiles of the distribution of interest so that scientists can capture specimens at varied locations in the environment based on differing values of some phenomenon of interest to them. For instance, if a marine biologist is interested in taking 9 physical water specimens that are spread over a range of chlorophyll concentrations, they may choose to use the deciles of the concentration. If only a small number of locations can be analyzed and the upper extrema values are of interest, they may choose to perform analysis at the locations of the (0.90, 0.95, 0.99) quantiles. The quantiles of interest

are a flexible way to describe the objective of specimen collection that is largely invariant to the exact phenomena being measured.

Our goal is to select locations for detailed analysis by scientists in two steps. We use a robot to perform an adaptive survey, and then find locations to suggest for specimen collection based on the measurements it takes. Specifically, we first aim to find the desired quantile values of the measurement distribution by adaptively selecting robot measurement locations that maximize an objective function designed to estimate quantile values, and then we produce suggested locations for specimen collection that are likely to contain these values. Figure 1 summarizes the approach, showing how our objective function for informative path planning and our loss function for location selection fit into an overall physical specimen collection pipeline.

Our contributions are:

- A planner with custom objective functions that adaptively improve the quality of estimated quantile values;
- A general loss function, which can be used with any optimizer, for selecting spatial locations for physical specimen collection which have the values estimated for the quantiles;
- Quantitative evaluation of our method on point and camera sensors with previously collected aquatic datasets;
- A demonstration of our method on a robot in a real-world crop health estimation task.

II. BACKGROUND

Online Informative Path Planning (IPP) consists of alternating between planning and taking an action, which usually corresponds to moving and measuring a value at the new location. A plan described by a partial trajectory p is created to maximize some objective function f over the measured locations \mathbf{X} and measured values \mathbf{Y} . The combined plan P is the concatenation of partial trajectories p . The plan and act steps iterate until the cost $c(P)$ exceeds some predefined budget [5], [6].

IPP as a Partially Observable Markov Decision Process (POMDP) provides a formulation for planning for taking measurements. POMDPs are a framework that can determine optimal actions when the environment is not fully observable or there is uncertainty in the environment. Formulating IPP as a partially observable Markov decision process (POMDP) has been explored [7]; previous works have used this for finding maxima [2] and to adapt the parameters of a rollout-based POMDP solver online to improve its efficiency [6]. We use these works to formulate and solve the IPP problem.

Using observations as measurements from the environment and a GP to represent the belief distribution, sequential Bayesian optimization can be formulated as a Bayesian search game [7]. Measurements from the environment $\mathbf{Y}_t = GT(\mathbf{X}_t)$ constitute observations which are partially observable components of the overall environment GT . To adapt the Bayesian search game formulation to IPP, the belief state is augmented with the state of the robot g_t and the actions the planner is allowed to take are restricted to local movements which

POMDP	Informative Path Planning
States	Robot position g_t , Underlying unknown function GT
Actions	Neighboring search points
Observations	Robot position g_t , Measured location(s) $\mathbf{X}_t = o(g_t)$ Measured value(s) $\mathbf{Y}_t = GT(\mathbf{X}_t)$
Belief	GP conditioned on previously measured locations $(\mathbf{X}_{0:t-1})$ and values $(\mathbf{Y}_{0:t-1})$
Rewards	$f(\mathbf{X}_t)$

TABLE I: **Informative Path Planning as a POMDP.** After [6], [7].

are feasible for the robot [2]. A complete description of the POMDP for IPP can be seen in Table I. We approximately solve this POMDP to select actions to take measurements for maximizing the objective function.

Gaussian Processes (GPs) are non-parametric models with uncertainty quantification which are widely used for IPP [2], [3], [6]. They approximate an unknown function from its known outputs by computing the similarity between points from a kernel function, in our case the squared exponential kernel [8]. GPs have been used to represent the belief distribution from observations in POMDP formulations of sequential Bayesian optimization and IPP [2], [6], [7]. In this work, we use a GP as the environment model to estimate the value $\mu(x)$ and variance $\sigma^2(x)$ at a specific location x .

Objective Functions for IPP are used to specify the mission objective to maximize during planning. We describe some history here and show that there is a gap in the current literature for estimating quantile values.

Often, the objective in IPP is to generate good coverage of a distributed phenomenon using an information theoretic objective such as entropy [3] or mutual information [4]. Our work differs from these by instead focusing on improving estimates for specific quantiles instead of broad understanding of the spatial makeup of a distribution.

Other common objectives include finding the highest concentration location [2] or measuring near hotspots [9]. Sequential Bayesian optimization objectives are used to locate extreme values or areas of high concentration and have shown success in doing so in simulation and in real field scenarios [2], [10], [11]. These specialized objective functions are typically tuned to locate a single point, often the extrema. There has also been work on extending IPP missions to domains such as underwater inspection by modifying the GP and objective functions [12]. Specialized objective functions have also been developed for the task of modeling continuous and discrete variables, as well as including sensor-specific characteristics such as the increase in uncertainty due to camera distance to the subject [13].

Other objective functions have been proposed for level set estimation which seek to classify points into sets above and below a pre-specified concentration, or a fraction of the concentration [14]. While these works have similarities with quantile estimation tasks, they require specifying a single quantile or are restricted to finding the extrema, as opposed to a set of arbitrary quantiles.

Physical Specimen Collection is the process of collecting portions of the environment for later analysis, which we choose using the quantiles of the distribution of the measured value in this work. The choice of physical specimen collection

locations is driven by the spatial heterogeneity of the studied phenomena [1], which include freshwater algal growth and crop health, two potential applications for this work.

We are primarily motivated by the study of algal growth and distribution in freshwater and marine ecosystems. Spatial heterogeneity of algal and cyanobacterial blooms is often obvious as conspicuous accumulations at down-wind and down-current locations at the surface of lakes, and vertical heterogeneity can occur due to water column stratification, differential growth at different depths, or active vertical migration by some planktonic species [15], [16]. Such heterogeneity thwarts studies designed to quantify the spatial distribution of algal biomass, and the concentrations of algal toxins that may be produced by some algae and cyanobacteria. Nonetheless, accurately characterizing this heterogeneity is fundamental for investigating average conditions (indicated by median quantiles) for investigating trends across many lakes or large geographical areas [17]. Assessing heterogeneity is also essential for assessing the worst case scenarios for exposure of animals and humans to algal or cyanobacterial toxins (specifically the highest quantiles). Such information contributes to ecotoxicological studies used to develop thresholds that constitute significant exposure to these toxins [18]. Robotic approaches for quantifying algal biomass across quantiles (often using chlorophyll or chlorophyll fluorescence as a proxy for algal biomass) have become a mainstay for quantitatively documenting heterogeneity in natural ecosystems [19], [20].

Adaptive surveys for selecting collection locations have typically focused on finding high-concentration areas from which to measure, and perform the specimen collection onboard the robot or by a following robot [21], [22]. Our work differs from such works as we do not seek to only find maximal concentration areas, but seek to find a variety of concentration values specified by the quantiles.

Estimation of Quantiles and Quantile Standard Error has been proposed using statistical techniques. We estimate the quantile value from measurements using $\tilde{v} = x_{[h]} + (h - [h])(x_{[h]} - x_{[h-1]})$ where $h = (n - 1)q + 1$, n is the number of measurements, and q is the quantile (Equation 7 in [23]). Standard error of a quantile estimate is a measure of the uncertainty, and is typically used in the construction of confidence intervals of a quantile estimate. Given a probability density function p , the standard error of the q th quantile is $\sqrt{q(1-q)}/(\sqrt{np}(\tilde{v}))$. Typically p is not known, but can be estimated from samples, using a density estimator [24]. Here we use a Gaussian kernel density estimator when planning.

Continuous Non-Convex Gradient-Free Optimization is used to solve general black-box optimization problems. We use these methods to select the analysis locations by minimizing our proposed loss function once the IPP survey is completed (see section IV-B).

The cross-entropy (CE) method is a popular approach which works by maintaining an estimate about the distribution of good solutions, and iteratively updating the distribution parameters based on the quality of the solutions sampled at each step, where the quality is determined via some function of

the solution configuration [25]. More precisely, a prior Pr and a posterior Po set of samples are maintained. Po contains n configurations sampled from $\mathcal{N}(\vec{\mu}, \vec{\sigma})$. $\vec{\mu}$ and $\vec{\sigma}$ are computed from Pr , which is the best $\eta\%$ configurations from Po . This continues iteratively, minimizing the cross-entropy between the maintained and target distributions [25].

Simulated annealing (SA) [26] is an iterative algorithm inspired by statistical mechanics. SA optimizes an energy function similar to a loss function in other optimization schemes, as configurations with lower energy are preferred. SA begins with a temperature $T = T_{\max}$. At each iteration, T is decreased exponentially, and the configuration is slightly perturbed randomly. The perturbed state is either accepted or rejected probabilistically based on its energy and T .

Bayesian optimization (BO) is a method for selecting parameters for difficult optimization problems such as hyperparameter selection [27]. BO is similar to IPP and sequential Bayesian optimization [2] in that these methods build a model using a GP and select points which maximally improve this model. BO uses acquisition functions, such as expected improvement, and iteratively selects the best point to sample from [28], [29].

III. FORMULATION

We use a grid-based representation of the planning space, $\mathbf{G}^\#$, which defines the set of locations that the robot could visit. For a robot that moves in \mathbb{R}^d , $\mathbf{G}^\# \subset \mathbb{R}^d$. $\mathbf{X}^\#$ is the set of locations the robot could measure. If the robot sensor has finer resolution than $\mathbf{G}^\#$, e.g. if the robot uses a camera sensor or takes measurements while traversing between grid points, then $|\mathbf{X}^\#| > |\mathbf{G}^\#|$. We define $\mathbf{X}_{0:t}$ as the locations the robot has measured up to time t and $\mathbf{Y}^\#$ and $\mathbf{Y}_{0:t}$ as the values at all possible measured locations, and the values the robot has measured up to time t , respectively.

We define the ground truth quantile values as $V = \text{quantiles}(\mathbf{Y}^\#, Q)$ where quantiles is the function described in Section II which computes the values V of the quantiles Q of a set of measurements, in this case $\mathbf{Y}^\#$. To define the robot's estimated quantile values, we compute $\tilde{V} = \text{quantiles}(\mu(\mathbf{X}^\#), Q)$ that is, the quantile values of the predicted values from the robot's current GP for all locations the robot could sense. This is done to prevent the number of measurements from which the quantile values are estimated from changing as the robot explores (instead of, e.g. using $\mu(\mathbf{X})$). By doing this, we ensure we always estimate the quantiles across the entire measurable area. During planning, we aim to minimize this error by taking actions which maximize an objective function f that minimizes the error in the quantile value estimate.

To suggest locations for the quantile values, we aim to find a set of $|Q|$ locations \mathcal{Q} in the continuous space whose values at those locations are equal to the quantile ground-truth values. A set of locations is defined as $\mathcal{Q} \in \mathcal{Q}^\#$ where $\mathcal{Q}^\# \subset \mathbb{R}^{d \times |Q|}$ and $\mathcal{Q}^\#$ is continuous over the space of $\mathbf{X}^\#$. In practice, the robot only has access to \tilde{V} (not V) during the selection process, so the problem of finding the estimated

¹ $\lfloor x \rfloor = \max\{a \in \mathbb{Z} | a \leq x\}$ and $\lceil x \rceil = \min\{b \in \mathbb{Z} | b \geq x\}$

quantile spatial locations \tilde{Q} can be stated as shown in eq. (1) with some selection loss function l_s (see eq. (7)).

$$\tilde{Q} = \arg \min_{Q' \subset Q^\#} l_s(\tilde{V}, Q') \quad (1)$$

IV. APPROACH

Figure 1 illustrates our method. We separate our approach into two steps: the survey, and the suggestion of locations. When performing the survey using IPP (Section IV-A), the robot takes measurements of the environment to improve its estimate of the quantiles. After the survey has concluded, location selection (Section IV-B) produces locations for scientists to visit to perform specimen collection.

A. Informative Path Planning

To plan which locations to measure, the robot uses a POMDP formulation of IPP. In order to generate a policy, we use the partially observable Monte Carlo planner (POMCP) [30]. POMCP uses Monte Carlo tree search to create a policy tree. To expand the tree and estimate rewards, the tree is traversed until a leaf node is reached. From the leaf node, the reward conditioned on that action is estimated by a random policy rollout which is executed until the discounted reward is smaller than some value ϵ . We modify the rollout reward to be fixed horizon, giving a reward of zero after a certain number of random policy steps. We adopt the t-test heuristic for taking multiple steps from a POMCP plan for IPP to improve performance of the planner with fewer rollouts [6]. We define $GP_{i-1} = GP(\mathbf{X}_{0:i-1}, \mathbf{Y}_{0:i-1}; \theta)$ as a GP conditioned on the previous locations ($\mathbf{X}_{0:i-1}$) and measurements ($\mathbf{Y}_{0:i-1}$) before measuring a proposed value, and $GP_i = GP(\mathbf{X}_{0:i-1} \cup \mathbf{X}_i, \mathbf{Y}_{0:i-1} \cup \mathbf{Y}_i; \theta)$ as a GP conditioned on the previous and proposed locations (\mathbf{X}_i) and measurements (\mathbf{Y}_i), where θ are GP parameters and $\mathbf{Y}_i = GP_{i-1}(\mathbf{X}_i)$. Because the observations $GT(x)$ for unseen locations are not known during planning, the predicted mean from GP_{i-1} is used [2].

Objective Functions for Quantile Estimation We develop two novel objective functions to improve the quality of quantile value estimates. Both compare a measure of the quality of the quantiles estimated by the GP before and after adding a measurement to the GP, and include an exploration term $c_{plan}\sigma^2(x_i)$ inspired by the upper confidence bound objective function [2], where c_{plan} is a chosen constant. For both proposed objective functions we use eq. (2), where δ is defined by the objective:

$$f(\mathbf{X}_i) = \frac{\delta(\mathbf{X}_i)}{|Q|} + \sum_{x_j \in \mathbf{X}_i} c_{plan}\sigma^2(x_j), \quad (2)$$

The first objective function, which we call *quantile change*, is based on the idea of seeking out values which change the estimate of the quantile values by directly comparing the estimated quantiles before and after adding the measured values to the GP. The idea behind this is that a measurement which changes the estimate of the quantiles indicates that the

quantiles are over- or under-estimated. This can be seen in in eq. (3):

$$\delta_{qc}(\mathbf{X}_i) = \|\text{quantile}(\mu_{GP_{i-1}}(\mathbf{X}^\#), Q) - \text{quantile}(\mu_{GP_i}(\mathbf{X}^\#), Q)\|_1 \quad (3)$$

The second objective function we develop, which we call *quantile standard error*, is based on the change in the estimate of the standard error for the estimated quantiles. It draws from the same idea that if the uncertainty in the quantile estimate changes after observing a measured value, then it will change the estimate of the quantile values, shown in eq. (4):

$$\delta_{se}(\mathbf{X}_i) = \|se(\mu_{GP_{i-1}}(\mathbf{X}^\#), Q) - se(\mu_{GP_i}(\mathbf{X}^\#), Q)\|_1 \quad (4)$$

se is an estimate of the standard error of the quantile estimate for quantiles Q . se uses a Gaussian kernel density estimate (we found it faster and more stable than other standard error estimators).

Baseline Objective Functions We compare against two baselines, one which maximizes spatial coverage of a phenomena and another which seeks maximal areas.

Entropy is a common objective function for IPP when only good spatial coverage of the environment is desired [3], [4], [31]. It provides a good baseline as it is often used when the specific values of the underlying concentration are unknown. Entropy is defined eq. (5):

$$f_{en}(\mathbf{X}_i) = \sum_{x_j \in \mathbf{X}_i} \frac{1}{2} \log(2\pi e \sigma^2(x_j)) \quad (5)$$

Another objective function we compare against is *expected improvement* (EI), which is widely used in Bayesian optimization and sequential Bayesian optimization for finding maxima [28], [29]. EI favors actions that offer the best improvement over the current maximal value, with an added exploration term ξ to encourage diverse exploration. The EI objective function is defined according to eq. (6):

$$f_{ei}(\mathbf{X}_i) = \sum_{x_j \in \mathbf{X}_i} I\Phi(Z) + \sigma(x_j)\phi(Z) \quad (6)$$

where $Z = \frac{I}{\sigma^2(x_j)}$, $I = \mu(x_j) - \max(\mu(\mathbf{X}^\#)) - \xi$, and Φ and ϕ are the CDF and PDF of the normal distribution, respectively.

B. Location Selection

Our final goal is to produce a set of $|Q|$ locations Q , at which the concentration values will be equal to V , the values of the quantiles Q . The selection process can be done offline as it does not affect planning. Finding locations that represent Q is difficult because the objective function over arbitrary phenomena in natural environments will likely be non-convex, and in a real-world deployment, the robot will only have an estimate \tilde{V} of the V it searches for.

With the location selection problem formulation as in Equation (1), we propose the loss function

$$l_s(\tilde{V}, \tilde{Q}) = \|\tilde{V} - \mu(\tilde{Q})\|_2 + c_{select}\sigma^2(\tilde{Q}), \quad (7)$$

where $c_{select}\sigma^2(\tilde{Q})$ is an added penalty for choosing points that the GP of collected measurements is not confident about², and l_s can be used in any optimization scheme. During optimization, eq. (7) is evaluated using \tilde{V} and returns the

²The parameters c_{plan} and c_{select} are distinct.

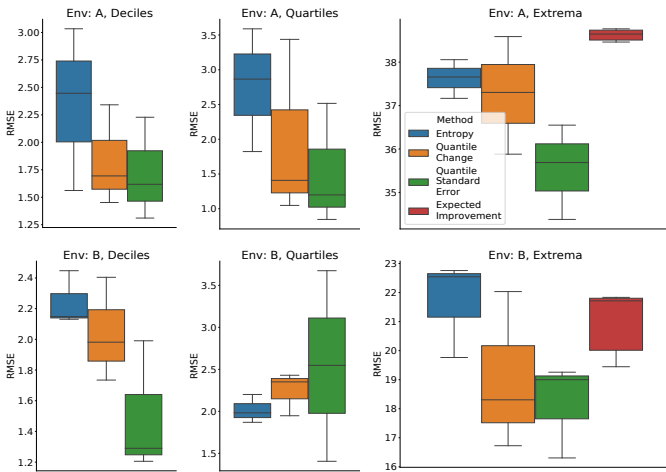


Fig. 2: **Simulated drone planning experiments with real data.** Error between ground truth quantile values and estimated quantile values. Datasets A and B collected using a hyperspectral camera in Clearlake, California. Units are 400nm channel pixel intensity (0 – 255). Each dataset and objective function pairing is run three times.

suggested specimen collection locations \tilde{Q} . We compare three optimization methods in our experiments to determine which best minimizes our selection loss function (eq. (7)). A strength of these types of optimization methods is that the formulation allows for suggesting points that may be spatially far from locations the robot was able to measure if they have values closer to the quantile values of interest.

V. EXPERIMENTS

To evaluate our approaches for planning and for location selection, we compare against baselines in two different informative path planning tasks in simulation using four datasets collected in the real world.

In the first task, a simulated drone with a virtual camera gathers data from orthomosaics (a single image produced by combining many smaller images, called orthophotos) collected of a lake using a hyperspectral sensor. The orthomosaics, A and B, are taken in the same location but on different days and times. The drone collects many measurements from one location, where each is a pixel in a downsampled image. As a proxy for chlorophyll concentration, we measure the 400nm channel pixel intensity (0 – 255). The drone maintains a constant altitude and moves in a 2D plane with a north-fixed yaw, moving in either the x or y direction per step.




Fig. 3: Experimental dataset distributions. Drone data is measured in pixel intensity; AUV data in $\mu\text{g}/\text{L}$ chlorophyll.




Fig. 4: **Simulated AUV planning experiments with real data.** Error between ground truth quantile values and estimated quantile values. Datasets C and D collected from a reservoir in California using an underwater robot with a chlorophyll fluorescence sensor. Units are $\mu\text{g}/\text{L}$ chlorophyll. Each dataset and objective function pairing is run three times.

In the second task, an autonomous underwater vehicle (AUV) explores an environment. Two AUV surveys, C and D, which were taken in the same reservoir but at different times and different areas, were conducted in a 3D lawnmower pattern using a chlorophyll sensor and are interpolated using a GP to $X^\#$. At each step, the AUV moves in one x , y , or z direction and takes five evenly spaced measurements when moving between locations.

We evaluate each task on their two respective datasets (A/B, C/D) and three different quantiles: deciles (0.1, 0.2, ..., 0.8, 0.9), quartiles (0.25, 0.5, 0.75), and upper extrema (0.9, 0.95, 0.99). See Figure 3 for a summary of the dataset distributions.

A. Informative Path Planning: Objective Functions

To evaluate how well our proposed IPP objective functions estimate the ground truth quantile values in real environments, we compare quantile change (eq. (3)) and quantile standard error (eq. (4)) against a baseline entropy objective function (eq. (5)). For the upper extrema quantiles, we also compare against expected improvement (eq. (6)), as it is similar to a sequential Bayesian optimization based IPP task.

1) *Setup*: In the planner, we use $\gamma = 0.9$, and each trial is run over 3 seeds. The objective c_{plan} parameter is




Fig. 5: Comparison between quantile standard error with POMDP solver and Spiral-STC coverage planner [32] on Drone Environment B estimating deciles, each over 3 trials.

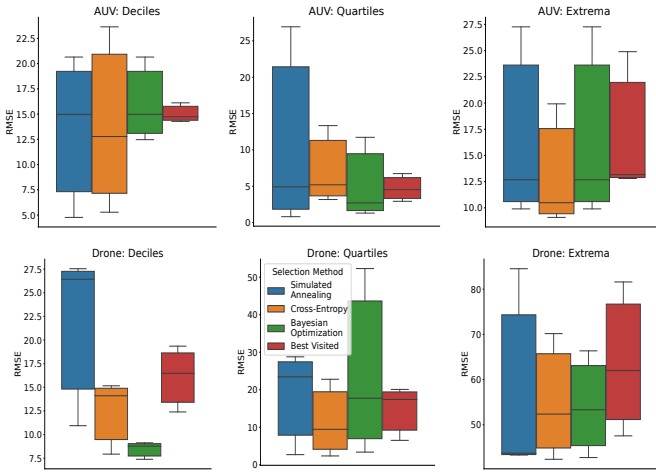


Fig. 6: **Quantile location selection results.** RMSE between the ground truth values at the selected locations and the corresponding true quantile values. Units are $\mu\text{g/L}$ chlorophyll for AUV, 400nm channel pixel intensity (0 – 255) for Drone.

set to the approximate magnitude of the rewards seen for each environment, which we found experimentally to be an adequate value. The GP mean was set to 0 and the datasets were normalized, while the lengthscale was set to 12. The exploration constant c_{plan} was set to $1E-6$ for quantile change and $1E-2$ for quantile standard error. To compare the performance of the approaches, we use the RMSE between the ground truth quantiles, V , and estimated quantile values after performing a survey, \tilde{V} .

Drone with Camera The drone is allowed to take 30 simulated pictures out of a grid with about 300 positions. Each picture is downsampled to 8×5 pixels (40 measurements) with 37.1° by 27.6° field of view, similar to the drone used in the field trial reported in section V-C. While this is downsampled from the true image, it captures the coarse trends of the concentration, which is both still scientifically useful and performant for GP evaluations. For each trial, the GP is seeded with 100 evenly spaced measurements across the workspace, as prior data. The planner uses 300 rollouts per step, and the maximum planning depth is 7.

AUV with Chlorophyll Sensor The AUV is simulated for 200 steps in a $12 \times 14 \times 2$ grid. The planner uses 130 rollouts per step and a maximum depth of 10. The GP is seeded with measurements from 50 locations.

2) *Discussion:* Overall, our results are robust across multiple experimental environments as well as robot sensor types. We find that planning with the quantile standard error objective function has a 10.2% mean reduction in median error across all environments when compared to using entropy. This shows that an objective function tailored to estimating a set of quantiles will outperform a coverage planner, such as entropy, that would typically achieve low overall error in reconstructing the entire environment model. The focus of our approach is to estimate the quantiles in a targeted way rather than to achieve uniform coverage.

Figures 2 and 4 show the results of planning with the proposed objective functions for the drone and AUV exper-

iments, respectively. In the drone tasks (A/B), quantile change and quantile standard error outperform the baseline entropy in estimating the deciles and upper extrema. In quartiles in environment A, both proposed methods perform well, but in environment B entropy outperforms our methods. For the extrema, both methods perform better than expected improvement, and entropy outperforms expected improvement in environment A. We believe this is because expected improvement focuses explicitly on improving the (single) maximal value and does not do a good job of localizing high concentration areas, thus overestimating the quantile values.

For the AUV tasks (C/D), quantile standard error outperforms entropy in environment C when estimating deciles and quartiles, and performs equally well as entropy in all other tasks besides estimating the extrema in environment D. Expected improvement performs poorly in estimating the extrema due to similar issues as with the drone. Quantile change performs poorly in most AUV tasks in contrast to the drone tasks, where it performs comparably. This may indicate that quantile change performance decreases with point sensors, or it is more sensitive to variation in the absolute measurement values. Thus, we recommend using quantile standard error unless there are computational constraints, since its performance is more consistent.

We also compare our approach using the proposed quantile standard error objective function to the Spiral-STC coverage planner [32]. Figure 5 shows that the proposed adaptive solution is able to reduce the error between the estimated quantiles and the ground truth quantiles more rapidly than the non-adaptive coverage planner.

B. Location Selection

To evaluate how close the values at locations suggested by our proposed loss function are to the estimated quantile values, we use the optimization algorithms simulated annealing, cross-entropy, and Bayesian optimization with the results from the quantile estimation task (Section V-A) using the quantile standard error objective function.

1) *Setup:* We compute the error by $RMSE(V, \tilde{Q}) = \|(V - GT(\tilde{Q}))\|_2$. We set c_{select} to 15 for deciles, 200 for quartiles, and 30 for extrema. We compare the results against a Best Visited (BV) baseline which selects the location the robot took a measurement at that is closest to each quantile value, i.e. solve $\tilde{Q}^* = \operatorname{argmin}_{\tilde{Q}_{bv} \subset \mathcal{X}} \|\tilde{V} - \mu(\tilde{Q}_{bv})\|_2$.

For SA, we use $T_{max} = 5$, $T_{min} = 0.001$, and cooling rate $cr = 0.005$ which leads to approximately 1000 optimization steps, and reset to the best solution every 100 steps. We start the optimization using the solution found by the BV baseline. For CE, we use $\alpha = 0.9$, $\eta = 0.9$, 50 samples per iteration, and 100 iterations. α is a weighting factor on new samples, which is used to prevent premature convergence. For BO, we use the expected improvement acquisition function and initialize the GP with 50 randomly selected Q_s as well as the solution found by the Best Visited baseline. We use 100 iterations and report the best found solution.

2) *Discussion:* The error between the values at the estimated quantile locations and the true quantile values for the




Fig. 7: **Physical locations (stars) selected by the cross-entropy optimizer for deciles on the drone experiment.** [Left] Black lines: true quantile value contours, overlaid on the absolute error between $\mu(X^\#)$ and $GT(X^\#)$. Note that the lowest error tends to follow the quantile contour lines. [Right] Red crosses: locations the robot visited, overlaid on the ground truth image.

drone and the AUV experiments can be seen in Figure 6. Overall, we find that cross-entropy and Bayesian optimization produced the locations with values closest to the true quantile values. Both these methods perform a global search in the space of possible locations, indicating that global search optimizes eq. (7) more effectively. Simulated annealing had greater variability in performance. We believe this is because it is a local search method and may fail to escape local optima. Best Visited produced a good initialization for the other methods, but was easily outperformed. We find that CE performs the best in three out of six possible scenarios. In particular, we see a 15.7% mean reduction in median error using CE with our proposed loss function compared to the BV baseline across all environments when using our proposed quantile standard error optimization function for quantile estimation during exploration.

In general, methods can find the best points when selecting for quartiles or deciles, while the upper extrema are more difficult. Because the robot is limited in the amount of environment it can explore, the upper extrema are less likely to be measured during exploration. This leads to these quantiles being more challenging to select representative points for.

Figure 7 shows results for one seed of the drone task when monitoring deciles. The suggested locations, shown as stars, align relatively closely with the true quantile values V , shown by the contours on the left image. This demonstrates the ability of the optimizers to produce good location suggestions to guide environmental analysis.

The right part of Figure 7 shows the same locations on top of the orthomosaic of what the drone could measure during exploration. This part of the figure highlights the difficulty of the problem of IPP for quantiles. The robot could only explore 15% of the total environment. With only partial knowledge of the distribution, the robot’s model of the phenomenon will vary based on the particular points it visited, which in turn affects the estimates of the quantiles. Similarly, Figure 8 shows the suggested locations for the upper extrema for one seed of the AUV task. For both tasks, all the optimization methods, with the exception of Best Visited, suggest points that may be spatially far from locations the robot has been able to visit. This allows for points with values potentially closer to the true quantile values to be selected for scientific analysis.




Fig. 8: **Physical locations (stars) selected by the cross-entropy optimizer for the upper extrema quantiles with an AUV using a chlorophyll point sensor.** Blue/green points are the measured locations X .

C. Field trial

The goal of our final experiment is to demonstrate our method on real hardware for a crop health monitoring task in an open, grassy field, where the objective is to estimate the deciles of the green present in each pixel of the images (we use green as a proxy for plant health).

1) *Setup*: We use a commercial, off-the-shelf drone with a standard camera to take measurements of the field at a constant height of 3m. Similar to the simulated drone, the drone in this experiment moves in a 2D plane with a north-fixed yaw. The drone has a limited budget of 20 pictures (planning steps) in a 16×16 m square grid, where $|G^\#| = 10 \times 10$ locations, and each picture is downsampled to 8×5 px. For planning, we use the quantile change objective function because it performs well for camera sensors and is faster than quantile standard error. We operate the drone with a custom app and use onboard GPS and IMU for localization.

2) *Discussion*: Figure 9 shows the resulting suggested locations using CE based on the 20 steps the robot took. We find that, although the robot cannot explore the entire workspace due to its limited budget, and in fact does not

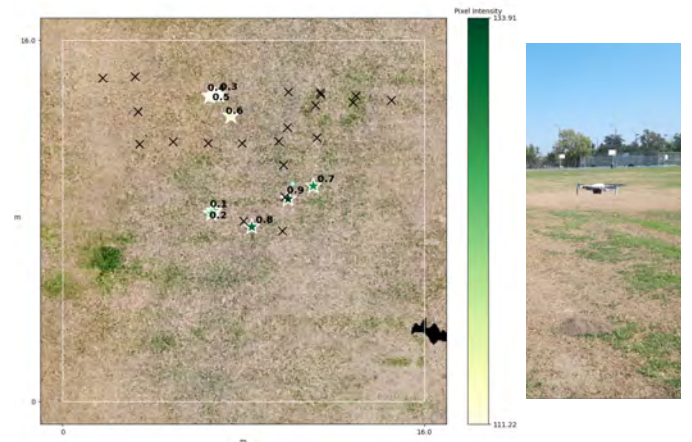


Fig. 9: **Visualization of a field trial modeling a crop health task.** *Left*: Crosses are locations where the drone took images. The drone is limited to only visit 20% of the possible locations to take images. The quantiles of interest are the deciles and the locations are chosen by cross-entropy. The 9 stars show locations suggested to collect physical specimen. The measurement of interest is the amount of green in each pixel. *Right*: The drone used in the field trial, in flight.

visit a large green area, the system is able to suggest varied locations for specimen collection. Some selected locations (e.g., those representing the 0.3, 0.4, 0.5 quantiles) are spatially close to each other, which suggests this area contains a large gradient, or those quantile values are similar. As the goal is to suggest locations that contain the estimated quantile values, such locations may be near one another.

VI. CONCLUSION

Scientists traditionally collect physical specimens at locations selected using heuristics. They later analyze these specimens in a laboratory to characterize a phenomenon of interest (e.g., the distribution of algae in the water). We propose to, instead, choose these specimen collection locations by first performing an informative path planning survey with a robot and then proposing locations which correspond to the quantiles of interest of the distribution.

To accomplish this, we propose two novel objective functions, quantile change and quantile standard error, for improving the estimates of quantile values through informative path planning. We test these in three settings: a drone with a camera sensor over lake imagery, an underwater vehicle taking chlorophyll measurements, and a field trial using a drone for a crop health task. Our objective functions perform well and outperform information-theoretic and Bayesian optimization baselines. In our experiments, our proposed quantile standard error objective function has a 10.2% mean reduction in median error when compared to entropy.

We also show that using our proposed loss function with black-box optimization methods can select environment locations for analysis that are representative of a set of quantiles of interests. We find that a cross-entropy optimizer using our loss function outperforms a baseline of using the best measured points, with a 15.7% mean reduction in median error in values across all environments.

Our approach can be used to guide physical specimen collection in real field scenarios, as demonstrated by our field trial. In recent work, we have discovered that some computational cost can be alleviated with a randomized sampling approach when evaluating the GP in the objective functions. We believe combining that and an approximate GP can improve scalability. Additionally, experimentally finding parameters was difficult for both the baselines and our proposed methodology, particularly for the GP; this is a direction for future work. We also plan to incorporate our specimen collection location suggestions into a larger field campaign involving multiple scientists.

REFERENCES

- [1] "High-resolution mapping of a harmful algal bloom: Understanding fine-scale environmental factors in clear lake, california," Tech. Rep., 2020. [Online]. Available: <https://video.ysi.com/webinar-expand-your-hab-horizons>.
- [2] R. Marchant, F. Ramos, and S. Sanner, "Sequential Bayesian optimization for spatial-temporal monitoring," in *UAI*, 2014.
- [3] S. Kemna, O. Kroemer, and G. S. Sukhatme, "Pilot Surveys for Adaptive Informative Sampling," in *ICRA*, 2018.
- [4] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in Gaussian processes," in *ICML*, 2005.
- [5] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *IJRR*, 2014.
- [6] C. E. Denniston, G. Salhotra, D. A. Caron, *et al.*, "Adaptive Sampling using POMDPs with Domain-Specific Considerations," in *ICRA*, 2020.
- [7] M. Toussaint, "The Bayesian Search Game," in *Theory and Principled Methods for the Design of Metaheuristics*, Y. Borenstein and A. Moraglio, Eds., 2014.
- [8] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT press, 2006.
- [9] S. McCammon and G. A. Hollinger, "Topological Hotspot Identification for Informative Path Planning with a Marine Robot," in *ICRA*, 2018.
- [10] A. Blanchard and T. Sapsis, "Informative Path Planning for Extreme Anomaly Detection in Environment Exploration and Monitoring," *arXiv:2005.10040*, 2021.
- [11] J. R. Souza, R. Marchant, L. Ott, *et al.*, "Bayesian optimisation for active perception and smooth navigation," in *ICRA*, 2014.
- [12] G. A. Hollinger, B. Englot, F. S. Hover, *et al.*, "Active planning for underwater inspection and the benefit of adaptivity," *IJRR*, 2013.
- [13] M. Popović, T. Vidal-Calleja, G. Hitz, *et al.*, "An informative path planning framework for UAV-based terrain monitoring," *Autonomous Robots*, vol. 44, 2020.
- [14] A. Gotovos, "Active learning for level set estimation," en, M.S. thesis, ETH Zurich, Zürich, 2013.
- [15] B. Seegers, J. Birch, R. Marin, *et al.*, "Subsurface seeding of surface harmful algal blooms observed through the integration of autonomous gliders, moored environmental sample processors, and satellite remote sensing in southern california: Harmful algal blooms subsurface seeding," *Limnology and Oceanography*, 2015.
- [16] A. Hozumi, I. Ostrovsky, A. Sukenik, *et al.*, "Turbulence regulation of microcystis surface scum formation and dispersion during a cyanobacteria bloom event," *Inland Waters*, 2020.
- [17] J. C. Ho, A. M. Michalak, and N. Pahlevan, "Widespread global increase in intense lake phytoplankton blooms since the 1980s," *Nature*, 2019.
- [18] A. C. Mehinto, J. Smith, E. Wenger, *et al.*, "Synthesis of ecotoxicological studies on cyanotoxins in freshwater habitats – Evaluating the basis for developing thresholds protective of aquatic life in the United States," *Science of The Total Environment*, 2021.
- [19] Y. Zhang, B. Kieft, B. W. Hobson, *et al.*, "Autonomous tracking and sampling of the deep chlorophyll maximum layer in an open-ocean eddy by a long-range autonomous underwater vehicle," *IEEE Journal of Oceanic Engineering*, 2020. DOI: 10.1109/JOE.2019.2920217.
- [20] S. L. Sharp, A. L. Forrest, K. Bouma-Gregson, *et al.*, "Quantifying Scales of Spatial Variability of Cyanobacteria in a Large, Eutrophic Lake Using Multiplatform Remote Sensing Tools," *Frontiers in Environmental Science*, 2021. DOI: 10.3389/fenvs.2021.612934.
- [21] J. Das, F. Py, J. B. Harvey, *et al.*, "Data-driven robotic sampling for marine ecosystem monitoring," *IJRR*, vol. 34, Oct. 2015.
- [22] S. Manjanna, A. Q. Li, R. N. Smith, *et al.*, "Heterogeneous multi-robot system for exploration and strategic water sampling," in *ICRA*, 2018.
- [23] R. J. Hyndman and Y. Fan, "Sample Quantiles in Statistical Packages," *The American Statistician*, vol. 50, 1996.
- [24] R. Wilcox, *Introduction to Robust Estimation and Hypothesis Testing (Third Edition)*, 3rd Edition, ser. Statistical Modeling and Decision Science. Academic Press, 2012, pp. 103–136.
- [25] P.-T. De Boer, D. P. Kroese, S. Mannor, *et al.*, "A tutorial on the cross-entropy method," *Annals of operations research*, vol. 134, 2005.
- [26] S. Kirkpatrick, C. D. Gelatt Jr, and M. P. Vecchi, "Optimization by Simulated Annealing," in *Readings in Computer Vision*, Elsevier, 1987.
- [27] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," in *Neurips*, vol. 25, 2012.
- [28] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient Global Optimization of Expensive Black-Box Functions," *Journal of Global Optimization*, 1998.
- [29] C. Qin, D. Klabjan, and D. Russo, "Improving the expected improvement algorithm," in *NeurIPS*, Long Beach, California, USA, 2017.
- [30] D. Silver and J. Veness, "Monte-Carlo Planning in Large POMDPs," in *NeurIPS*, vol. 23, 2010.
- [31] C. Denniston, A. Kumaraguru, and G. S. Sukhatme, "Comparison of Path Planning Approaches for Harmful Algal Bloom Monitoring," in *OCEANS MTS/IEEE*, 2019. DOI: 10.23919/OCEANS40490.2019.8962687.
- [32] Y. Gabriely and E. Rimon, "Spiral-stc: An on-line coverage algorithm of grid environments by a mobile robot," in *ICRA*, 2002.