2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 23-27, 2022, Kyoto, Japan

# Adaptive Sampling of Latent Phenomena using Heterogeneous Robot Teams (ASLaP-HR)

Matthew Malencia, Sandeep Manjanna, M. Ani Hsieh, George Pappas, Vijay Kumar

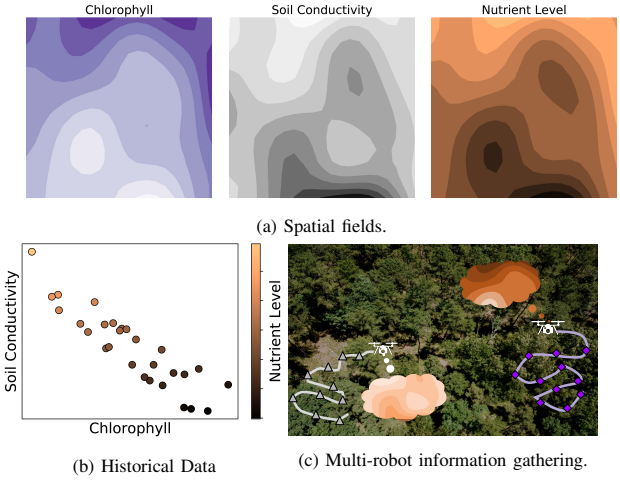
2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) | 978-1-6654-7927-1/22/$31.00 ©2022 IEEE | DOI: 10.1109/IROS47612.2022.9982270

***Abstract*—In this paper, we present an online adaptive planning strategy for a team of robots with heterogeneous sensors to sample from a latent spatial field using a learned model for decision making. Current robotic sampling methods seek to gather information about an observable spatial field. However, many applications, such as environmental monitoring and precision agriculture, involve phenomena that are not directly observable or are costly to measure, called latent phenomena. In our approach, we seek to reason about the latent phenomenon in real-time by effectively sampling the observable spatial fields using a team of robots with heterogeneous sensors, where each robot has a distinct sensor to measure a different observable field. The information gain is estimated using a learned model that maps from the observable spatial fields to the latent phenomenon. This model captures aleatoric uncertainty in the relationship to allow for information theoretic measures. Additionally, we explicitly consider the correlations *among* the observable spatial fields, capturing the relationship between sensor types whose observations are not independent. We show it is possible to learn these correlations, and investigate the impact of the learned correlation models on the performance of our sampling approach. Through our qualitative and quantitative results, we illustrate that empirically learned correlations improve the overall sampling efficiency of the team. We simulate our approach using a data set of sensor measurements collected on Lac Hertel, in Quebec, which we make publicly available.**

***Index Terms*—Informative Path Planning, Heterogeneity, Multi-Robot Systems**

## I. Introduction

Robot teams have proven to be powerful tools for environmental monitoring, drastically improving efficiency over fixed sensor networks or manual sampling [1]. This effectiveness is driven in large part due to their ability to move and collect samples from multiple locations, and due to the sampling strategies that decide where measurements should be taken [2]. However, many applications involve **phenomenon that either cannot be directly measured by sensors or require slow and costly sensors to measure, called *latent phenomena* or *latent spatial fields***. For example, in agricultural applications, chlorophyll is measured as a proxy for health or as an indicator of an infestation [3]. And in marine settings, scientists infer aquatic health by measuring various chemical concentration indicators [1]. In this paper, we answer the question: **how can robot teams choose sensing locations that best inform a latent phenomenon?** We formulate this question as an Adaptive Sampling of Latent Phenomena (ASLaP) problem and propose an informed path planning technique for a heterogeneous team of robots (ASLaP-HR) that seeks to estimate a quasi-static latent spatial field. We assume that the observable and latent spatial fields change slowly or remain static at least for the time span of a robotic survey and that robot team's communication is not constrained with the region.




(a) Spatial fields.

(b) Historical Data

(c) Multi-robot information gathering.

Fig. 1: An example scenario: *(a)* Farms can be described by various spatial fields. A spatial map of nutrient levels is most informative to a farmer, but it requires costly and time-consuming sampling and analysis of leaves. Chlorophyll and soil conductivity spatial fields can be measured quickly and cheaply by in-situ sensors. *(b)* Historical data is used to train a model to predict nutrient levels using chlorophyll and soil conductivity measurements. *(c)* A team of robots can measure chlorophyll and soil conductivity relatively efficiently and use the learned model from (b) to make real-time decisions about where and what to sample next to get best estimates of nutrition.

Consider the example scenario presented in Fig. 1: A farmer monitors plant nutrition by sending leaf samples to a lab for analysis. However, cheap and fast sensors, such as chlorophyll and soil conductivity sensors, can provide in-situ proxy measurements that estimate plant nutrition. Our approach to this scenario uses historical data of co-located chlorophyll, soil conductivity, and analyzed leaf samples to learn the relationship between the sensor measurements and nutrient levels. Then, a team of robots (equipped with soil conductivity and chlorophyll sensors) collaboratively estimates the nutrition levels of the whole farm (or plot) by using the learned model to inform where and which sensor measurements will provide the best estimate of nutrition.

Current work on sampling approaches for robot teams either optimize objectives on observable spatial fields [4], [5] or assume an exact proxy for a latent spatial field [6].

We gratefully acknowledge the support from ARL Grant DCIST CRA W911NF-17-2-0181, NSF Grant CNS-1521617, ARO Grant W911NF-13-1-0350, ONR Grants N00014-20-1-2822 and ONR grant N00014-20-S-B001, and Qualcomm Research. The first author acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1845298.

The authors are with GRASP Laboratory, Levine Hall 4th floor, University of Pennsylvania, 3330 Walnut Street, Philadelphia, PA 19104-6228. *corresponding authors emails:* {malencia,msandeep}@seas.upenn.edu



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

However, this does not necessarily align with optimizing relevant information of the latent phenomenon of interest [7]. After collecting data of their observable spatial fields, these teams rely on human experts for higher level decision making about the phenomenon of interest.

Heterogeneity in robot teams has been considered in robotic sampling [8], with a focus on diverse robot types, e.g., aerial vs ground robots [9]. While some papers consider sensor heterogeneity, this area is limited to sensors that align with prescribed roles [6] and sensors that are homogeneous in type but heterogeneous in cost, range, and accuracy [10]. In this work, we investigate **heterogeneity in sensor types, where each type measures a different spatial field**. These diverse observations can combat uncertainty of the latent phenomenon because real world sensor spaces are not independent and map nonuniformly to the latent spatial field.

In order to leverage the strength of heterogeneous sensors, we use data-driven methods to model the relationship between the observable spatial fields and the latent phenomenon. Most machine learning methods are limited by the resolution of the input data, e.g., using satellite data to make predictions about entire plots of land rather than individual locations or crops [11], or they postprocess spatial data to gain meaningful insights. Our work seeks to enable predictions that reason directly about the latent spatial field in real time while also providing the precision of studying localized phenomena.

### A. Related Work.

There are myriad problem spaces that involve path planning for mobile robots. Bayesian optimization approaches seek to find the maximum of a spatial field [12]–[14]; exploration and active mapping optimize the observation of an unknown area [15]–[17]; and active information gathering or target tracking focus on estimating the state of dynamic hidden states [18]–[20]. The objective functions studied among these applications can be categorized as information-theoretic driven [21], [22], such as mutual information [23], [24], and sample driven, as is the case of hotspot sampling [25], [26]. This work seeks to maximize the expected information gained by a team of heterogeneous robots.

Heterogeneity of multi-robot teams can vary drastically. Manjanna *et al.* use two autonomous surface vehicles, one with a water quality sensor and the other with a water sampler, where the sensor vehicle explores to inform the sampler's decisions [6]. Similarly, Manderson *et al.* use an aerial vehicle to inform the decisions of a ground vehicle [9]. Chand and Carnegie investigated heterogeneous processing and sensing capabilities to assign tasks to different robot types, where robots with low-value sensing do computation rather than exploration tasks [27]. Shi *et al.* use a Voronoi approach where the space is partitioned according to heterogeneous capabilities, like mobility constraints in ground versus aerial vehicles [10]. Similarly, Cai *et al.* consider heterogeneous capabilities such as field of view and traversability of different environment types [28]. Others still consider resource constraints such as energy budgets and fuel limitations [29]–[31]. Whereas Chand and Carnegie model different levels of sensing ability (e.g., low vs high), we model qualitatively different sensor types, which all take high quality sensor measurements in their respective observable space. Manjanna *et al.* and Manderson *et al.* investigate different robot types (sensor vs sampler, and aerial vs ground sensing, respectively), but their investigations focus on two prescribed roles. In this work, we do *not* consider heterogeneity in robot resources (e.g., computation power, or mobility type) but rather heterogeneity in sensor types; we model an arbitrary number of sensor types where tasks and roles are not pre-designed.

Recent work has only just begun to develop models and algorithms to make real-time decisions with respect to a latent phenomenon of interest. Arora *et al.* study multi-modal active perception, where measurements must be combined with domain knowledge to deduce information about variables or phenomena that cannot be directly measured [7]. Their paper investigates a single robot with multiple sensors that make discrete observations and maps the observations to the variables of interest through an expert-designed Bayesian network. In this work, we consider a multi-robot team where sensor measurements are continuous and we learn a continuous representation mapping from the observations to the latent phenomenon of interest.

### B. Contributions

We formalize adaptive informative sampling for multiple continuous sensors to maximize information theoretic objective of a latent phenomenon. We use a robot team with heterogeneous sensor types to measure observable spatial fields. A learned model from historical data enables the robot team to infer about the latent spatial field in real time and use these inferences to choose sample locations that maximize the information gain of the latent phenomenon. Additionally, we learn correlations **among** the multiple observable spatial fields; these correlations are learned offline and are used to make real-time sampling location decisions. We investigate the impact of these learned correlations on the adaptive sampling of the latent field. Our main contributions are:

- We formulate the ASLaP problem as the maximization of the expected information of a latent spatial field.
- We develop ASLaP-HR, an adaptive sampling approach for a team of robots with heterogeneous sensors using a learned mapping from observations to the latent phenomenon.
- We show that the learned correlations among the observable spatial fields improve the sampling efficiency of our adaptive sampling approach.

In Section II, we present a novel ASLaP problem formulation that is not restricted to myopic or non-myopic approaches, nor does it restrict motion uncertainties or arbitrary sets of actions,[1] though the methods presented in Section III are myopic and do not consider motion uncertainty. We validate our approach by running experiments on real sensor

[1]The problem formulation considers a discrete action set. Using continuous dynamics would require slight reformulation.



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

data collected on a lake. The sensor data, which we make publicly available [2], includes co-located measurements of ten water quality parameters. Our experiments use three of the ten sensor fields, two as the sensor spaces and one as the latent variable of interest. The experimental results show that the learned correlations among the observable spatial fields improve overall sampling. Future work will investigate the online learning of these correlations.

## II. Problem Definition

Consider a region of locations $x \in \mathbb{R}^2$ with a spatial field $\phi : \ x \mapsto \mathbb{R}$, such as temperature.[3] Adaptive informative sampling seeks to estimate this spatial field $\phi$. The estimate is a belief model, $\hat{\phi} : x \mapsto \mathbb{D}$, where $\mathbb{D}$ is the set of distributions in $\mathbb{R}$; for each $x$, $\hat{\phi}(x)$ is a one-dimensional distribution that estimates $\phi(x)$. To model $\hat{\phi}$, robots make observations at locations $x$ that maximizes information gain, $\mathcal{I}$. Often, $\mathcal{I}$ is calculated using Shannon entropy, $H$. The information gain with respect to $\phi$ of a measurement at $x$ is the change in entropy of the field estimate due to the new measurement: $\mathcal{I}_{\hat{\phi}}(x) = H(\hat{\phi}) - H(\hat{\phi}|x)$. Additionally, adaptive sampling problems often include resource constraints, such as energy budgets [32] or sampling budgets [6]. We limit the number of measurements using a time horizon constraint, $T$.[4] Therefore, a robot seeks a series of measurements $x_1, \ldots, x_T$ that maximize $\sum_{t=1}^{T} \mathcal{I}_{\hat{\phi}}(x_t)$. This extends to a team of $N$ robots where the vector $\vec{x}_t = (x_t^{(1)}, \ldots, x_T^{(N)})$. This type of problem formulation is found in the previous literature on informative path planning [7], [25], [33], [34]. Conversely, we consider a latent spatial field where $\phi$ **is not directly observeable**, yielding the Adaptive Sampling of Latent Phenomena (ASLaP) problem formulation:

**Problem 1.** ASLaP: *Given a team of $N$ agents, find observation locations $\mathbf{x}_t$ at all time steps in time horizon $T$ that maximize the sum of information gained relative to the latent spatial field $\phi$. Formally:*

$$\underset{\vec{x}_1,\ldots,\vec{x}_T}{arg\,max} \sum_{t=1}^{T} \mathcal{I}_{\hat{\phi}}\left[\vec{x}_t\right] \tag{1}$$

The novel problem formulation in Problem 1 is challenging because $\phi$ **is a latent phenomenon**. This breaks the fundamental assumption that the robots can directly measure the field $\phi$, thus requiring new methods that enable inference over this latent spatial field. In addition to this novel problem formulation, this work is the first to take advantage of learned spatial field correlation models to improve adaptive sampling.

In the following section, we present an approach using a team of robots with heterogeneous sensors, where the different sensor types measure observable spatial fields to collectively inform the estimate of the latent phenomenon by using a mapping that is learned from historical data. Additionally, we investigate how learned correlation models among these observable spatial fields can be leveraged to improve the efficiency of this team's sampling.

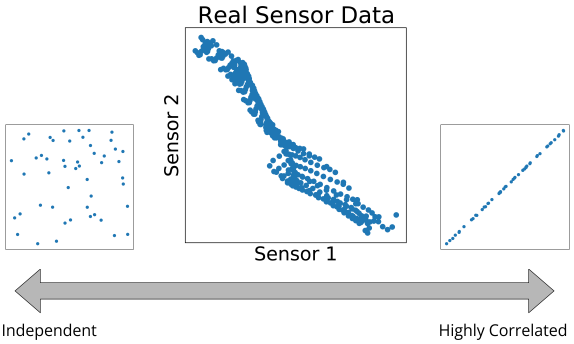



Fig. 2: The mutual information between different sensor types. Independent sensor types are independent sources of information. For perfectly correlated sensor types, co-located measurements from a second sensor type provide zero additional information. In reality, different sensor types (and the respective measurable spatial fields), have complex dependencies with mutual information that is complex, unknown function of the measurements. The middle scatter plot above shows the relationship of real sensor data that is used in our experiments in Section IV.

## III. Approach

We consider a heterogeneous team of $N$ robots with $S$ sensor types among the team, where $S = N$, which means that each robot has a distinct sensor type. It is possible to extend our approach to an arbitrary number of robots sensor combinations where $S \neq N$.[5] Each sensor type measures a spatial field, and, as with real world observeable spatial fields, sensor measurements from different sensor types are *not* independent of one another, as seen in Fig. 2. We use the index $s = 1, \ldots, S$ for both robots and sensors. Each robot $r_s$ can measure a spatial field $z^{(s)} : x \mapsto \mathbb{R}$. Let $x_t^{(s)}$ be the location of robot $r_s$ at time $t$ and let the respective observation be $z^{(s)}(x_t^{(s)})$. For ease of notation, the observation made by robot $r_s$ at time $t$ will be denoted by $z_t^{(s)}(x)$. Let $Z : x \mapsto \mathbb{R}^S$ be the $S$-dimensional observable spatial field, where $Z(\vec{x}) = \left[z^{(1)}(x), \ldots, z^{(S)}(x)\right]$. The spatial field $Z$ is unknown but observable, therefore, online measurements are used to maintain a belief estimate $\hat{Z} : x \mapsto \mathbb{D}^S$.

With this team of heterogeneous multisensor robots, we seek to build an estimate of the latent spatial field. To do this, a mapping $g$ from observations in $Z$ to the latent spatial field $\phi$ is learned from historical data. Robots make measurements of $Z$, build an estimate of the observable spatial field $\hat{Z}$, and then use the learned mapping $g : \hat{Z} \mapsto \hat{\phi}$ to maintain the estimate of the latent spatial field $\phi$. Learning $g$ is challenging because the mapping from $\hat{Z} \in \mathbb{D}^S$ to $\hat{\phi} \in \mathbb{D}$ has aleatoric uncertainty, meaning even an exact measurement of an observable spatial field yields an uncertain latent spatial field value. Multiple machine learning methods can capture this aleatoric uncertainty [35], such as

[2]This data is available at: https://github.com/SandeepManjanna/water_quality_dataset

[3]We study 2-dimensional spatial fields in this work, though 3-dimensional space could be considered.

[4]The constraint $T$ can be interpreted as an energy budget, assuming each action takes equal energy, or as a sensing budget of at most $T$ measurements per robot (we assume one measurement is taken each time step). Other works define a cost per action/measurement and constrain the sum of costs.

[5]For $S \neq N$, there are $N^{(s)}$ robots per sensor type such that $\sum_{s=1}^{S} N^{(s)} = N$.



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

the Gaussian Process approach that we use in Section IV. Because the goal of our approach is to create a good estimate of the latent phenomenon, the observation locations $x$ are chosen to maximize the information gain with respect to $\phi$. Using the learned mapping $g$, the information gain is now calculated by $\mathcal{I}_{g(\hat{Z})}(x) = H(g(\hat{Z})) - H(g(\hat{Z}|x))$, which yields the Adaptive Sampling of Latent Phenomena using Heterogeneous Robots (ASLaP-HR) approach:

**ASLaP-HR.** *Given a team of $S$ robots, each equipped with a distinct sensor $s$, each robot $s$ takes measurements of the field $z^{(s)}$ to collectively form an estimate $\hat{Z}$ of the entire observable spatial field. This in turn allows estimates of the latent spatial field by using $g$, a mapping from the observable spatial fields to the latent spatial field that must be learned from historic data. At each time step $t$, choose the observation location for each robot, $\vec{x}_t = (x_t^{(1)}, \ldots, x_T^{(N)})$, that maximize the information gain $\mathcal{I}_{g(\hat{Z})}$ with respect to this latent phenomenon. Formally:*

$$\underset{\vec{x}_1,\ldots,\vec{x}_T}{arg\,max} \sum_{t=1}^{T} \mathcal{I}_{g(\hat{Z})}\left[\vec{x}_t\right] \tag{2}$$

While prior works reason about observable spatial fields, ASLaP-HR reasons about a latent phenomenon through $g$. The following section presents an ASLaP-HR planning algorithm, which leverages a learned correlation among the observable spatial fields for improved information gathering. Section IV presents a specific implementation of this algorithm and shows the improvement gained by leveraging the learned correlation.

### A. Planning Algorithm Details

We present Algorithm 1 for ASLaP-HR, where a team of $S$ robots take measurements of $Z$, build an estimate of the observable spatial field $\hat{Z}$, and then use the learned mapping $g$ to build an estimate of the latent spatial field $\hat{\phi}$. Algorithm 1 takes the following inputs:

- $\vec{x}_0 \in \mathbb{R}^{2xS}$ initial positions of all robots
- $T \in \mathbb{Z}^+$ time horizon
- $g : \hat{Z} \mapsto \hat{\phi}$ mapping from observations to latent field
- $h_s : z^{(s)} \mapsto Z$ correlation model for sensor type $s$

Algorithm 1 begins with each robot $r_s$ observing the spatial field $z^{(s)}$ at their initial position $x_t^{(s)}$, yielding measurements $Z(\vec{x}_0)$. These measurements are used to update the belief $\hat{Z}$ (line 3) which is then used by $g$ in line 4 to update the belief $\hat{\phi}$. In Section IV, we model $\hat{Z}$ using Gaussian Processes, which produce inputs to the learned model $g$, which in turn outputs $\hat{\phi}$. The team then communicates the sensor measurements.[6] For each time step $t$, and robot $r_s$, the next waypoint $x_t^{(s)}$ is selected to maximize the expected information gain. Calculating the information gain using Shannon Entropy (as shown in Section II) is computationally intractable, therefore, the variance of $\hat{\phi}$ is used as an estimate of the expected information gain [36].

Next, $sense(s, x_t^{(s)})$ on line 8 returns a new sensor measurement $z_t^{(s)}$. This new sensor measurement is used to refine the model $h_s$ using any measurements from other sensors that are located together at $x_t^{(s)}$. We assume that GPS localizes the robots; the uncertainty in this localization does not hinder the estimate of the beliefs. The new sensor measurements are communicated, and the correlation models $h_s$ are used in line 11 to update the beliefs $\hat{Z}$ at the points $x_t^{(s)}$. Line 12 updates the belief $\hat{Z}$, and line 13 updates the belief $\hat{\phi}$. Algorithm 1 then repeats this sequence to iterate over all $T$ time steps, at the end returning the most up-to-date belief $\hat{\phi}$ of the latent phenomenon. The next section elaborates on the line 11 update of $\hat{Z}$.

**Algorithm 1** Pseudocode for ASLaP-HR Planner

**Input:** $\vec{x}_0, T, g, h_1, \ldots, h_S$
**Output:** Estimated latent spatial field $\hat{\phi}$

1: Initialize $\hat{\phi}, \hat{Z}$
2: Measure $Z(\vec{x}_0)$
3: $\hat{Z} \leftarrow update\left(Z(\vec{x}_0)\right)$
4: $\hat{\phi} \leftarrow g\left(\hat{Z}\right)$
5: **for** $t = 1, \ldots, T$ **do**
6:  **for** $s = 1, \ldots, S$ **do**
7:  $x_t^{(s)} \leftarrow \underset{x}{arg\,max} \quad \mathcal{I}_{g(\hat{Z})}\left[\vec{x}_t\right]$
8:  $z_t^{(s)} = sense(s, x_t^{(s)})$
9:  $h_s \leftarrow refine(x_t^{(s)}, z_t^{(s)})$
10:  **end for**
11:  $\hat{Z}(x_t^{(s)}) = h_s\left(z_t^{(s)}\right) \;\forall\, s = 1, \ldots, S$
12:  $\hat{Z} \leftarrow update\left(Z(\vec{x}_t)\right)$
13:  $\hat{\phi} \leftarrow g\left(\hat{Z}\right)$
14: **end for**
15: **return**

### B. Learned Correlations among the Observable Spatial Fields

Through this point in Section III, we have presented the basic structure of Algorithm 1. The last core piece, presented here and investigated in Section IV, is the learned correlations among the observable spatial fields $h_s$. Algorithm 1 does not require line 11. However, we hypothesize that **heterogeneous robots can leverage the relationships between sensor-types to achieve more efficient information gathering.**

When $g$ updates the belief of the latent spatial field at a point $x$, it calculates $\hat{\phi}(x)$ as a function of the entire observable field $\hat{Z}(x)$. Therefore, the precision of the value and variance of $\hat{\phi}(x)$ depends on the precision of the entire belief $\hat{Z}$. However, a robot $r_s$ measures only part of $\hat{Z}$, specifically $z^{(s)}(x)$.

The precision of the variance of $\hat{\phi}(x)$ is especially important because it informs the selection sensing locations. If $\hat{\phi}(x)$ *underestimates* variance, line 7 calculates a lower expected information gain than is actually present, causing robots

[6]We assume robots can communicate with each other within the region, and that time needed for communication and belief updating is significantly shorter than the time-step duration.



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

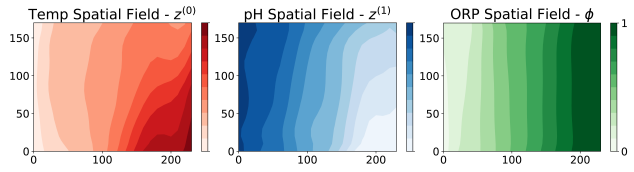



Fig. 3: Our simulated environment includes two observable spatial fields: $z^{(0)}$ temperature (left), and $z^{(1)}$ pH (middle). These fields are unknown but observable:f $z^{(0)}$ can be measured by robot $r_0$ and $z^{(1)}$ can be measured by robot $r_1$. The simulated environment also includes a latent spatial field $\phi$ of ORP values (right) which **cannot be measured by either robot.**

to avoid informative measurements. If $\hat{\phi}(x)$ *overestimates* variance, robots may take redundant measurements because they expect more information than is actually available.

In Algorithm 1, we see that the observable spatial field correlation models $h_s$ are used to improve the estimate of $\hat{Z}$, which in turn improves the estimate of $\hat{\phi}$. Given a measurement $z^{(s)}(x)$, the correlation model $h_s$ estimates the value (with variance) of the entire observable spatial field $\hat{Z}(x)$. These models capture the aleatoric uncertainty in the correlations among the observable spatial fields. For example, if there is low aleatoric uncertainty between two sensor types, one sensor acts almost as a noisy sensor of another spatial field. Conversely, correlations with high aleatoric uncertainty only marginally diminish the variance of other spatial fields. The line 11 of Algorithm 1 shows how the correlation models are used after a robot moves to a new $x$ and takes a new sample $z^{(s)}$. The new sample is the input to model $h_s$ that provides updated estimates of the value and variance of the entire observable space at that location, $\hat{Z}(x)$.

## IV. Results

We present an instantiation of Algorithm 1 and investigate the impact of the correlations among the observable spatial fields described in Section III-B. This section presents simulated experiments where a team of robots estimate a latent phenomenon. We use real sensor data collected from a lake to create a simulated environment of two observable spatial fields and a latent phenomenon of interest. In the following subsections, we present a marine data set, we overview a simple myopic planner that instantiates Algorithm 1, we describe the modeling of the correlations between the observable spatial fields, and we present results from simulated experiments. We show qualitatively that the learned correlations between the observable spatial fields $h_s$ improve the performance by causing the robot team to spread their samples out. Our quantitative evaluations show that these correlations improve the accuracy of the latent field estimate.

### A. Experimental Setup - Data and Simulated Environment

We simulate a team of two robots in a marine setting. One of the robots is equipped with a temperature sensor and the other one with a pH sensor. Together, they are seeking information about oxidation reduction potential (ORP), the latent phenomenon of interest in these simulated experiments. Historical data temperature, pH, and ORP are needed to learn the correlation model $g$.

*Lac Hertel Data.* The historical data used in our simulated experiments was collected on Lac Hertel in Mont Saint-Hilaire, Quebec. The full data set, detailed in the Appendix, is made publicly available. For this paper, we use temperature and pH data as our observable spatial fields, and the oxidation reduction potential (ORP) data as the latent phenomenon. ORP measures the ability of a lake or river to cleanse itself or break down waste products; ORP is aligned with the "knowledge" that might be sought in the environmental monitoring of a lake because it indicates whether drinking water is sanitary or if the water is suitability for anaerobic microbial processes. Additionally, we found that ORP, pH, and temperature had well-formed relations that could be modeled. Let the historical data $\mathcal{D} = (\mathcal{X}_\mathcal{D}, d_0, d_1, d_2)$ be a tuple of the sets of sample locations, $\mathcal{X}_\mathcal{D}$, the sets of measurements from two observable spatial fields (temperature:$d_0$ and pH:$d_1$) and the set of measurements from the latent spatial field (ORP:$d_2$). This historical data is used to learn the mapping from observations to the latent spatial field $g$ and the correlations $h_0$ and $h_1$.

*Simulated Environment.* To test Algorithm 1, we create a simulated environment of observable and latent spatial fields. Fig. 3 shows the three spatial fields of our simulated test environment. All three fields are unknown. However, $z^(0)$ can be measured by robot $r_0$ and $z^(1)$ can be measured by robot $r_1$. The latent field $\phi$ cannot be measured by the robot team; this spatial field is used only to evaluate the performance of our algorithm. We normalize all spatial fields to a scale from zero to one.[7] To simulate our robots' motion in this simulated environment, we discretize the region into a grid, defining the set $\mathcal{X}$ of locations and corresponding ground truth sensor measurements $\mathcal{GT} = (gt_0, gt_1, gt_2)$.

### B. Myopic Planner

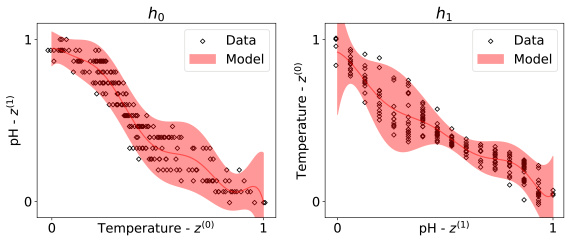
For our simulated experiments, we assume a robot can move to any adjacent cell during a time step. For a given location $x_{t-1}^{(s)} \in \mathcal{X}$, the possible next locations $x_t^{(s)}$ in line 7 are the eight grid cells surrounding the current location (excluding cells outside of the boundary, i.e., $\notin \mathcal{X}$). Lines 3 and 12 create belief models of the observable spatial field $\hat{Z}$. At each time step $t$, the sensor measurements and respective sample locations are used to fit the belief model $\hat{Z}$ of the observable fields using Gaussian Processes. To create these beliefs, we assume known spatial correlations (covariance kernels), which is a common assumption because underlying physics informs these correlation (e.g., how temperature changes spatially). The belief model of the latent phenomenon, $\hat{\phi}$, is created/updated using the pre-learned mapping $g$ to predict the mean and variance of the latent spatial field. The input to $g$ is $\hat{Z}$ and the output is $\hat{\phi}$. Importantly, we do not assume known spatial correlations of the latent variable; Mean and variance predictions are strictly a function of the observable spatial field belief model and the mapping $g$.

The mapping $g$ can be learned by various machine learning methods which capture the aleatoric uncertainty between

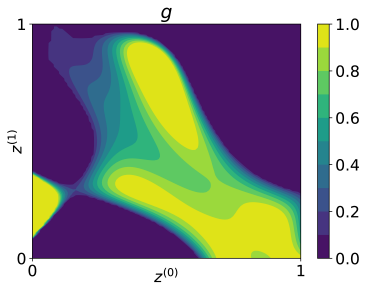
[7]Normalizing the sensor data aids in learning correlation models.



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.




(a) Correlation models




(b) Mapping $g$ from $Z$ to $\phi$

Fig. 4: *(a)* The learned correlation mappings $h_0$ from the pH spatial field to the temperature spatial field and $h_1$ from temperature to pH. The solid line shows the mean prediction and the bars show the variance prediction. *(b)* The learned mapping $g$ from observations to the latent spatial field.

observations and the latent spatial field [35]. Because we are only considering $S = 2$ observable spatial fields, deep learning methods are not required. We use Gaussian Process Regression with inputs from $d_0$ and $d_1$ and targets of $d_2$ to train our mapping $g$. Then, to capture the uncertainty in the latent variable, we use samples at each location of our belief $\hat{Z}$ to capture the variance in the belief; we predict the value and variance of the latent spatial field for all samples; then, we calculate the mean and variance of the latent spatial field predictions at each point to create our latent belief $\hat{\phi}$.

To calculate the expected information gain in line 7, we use the variance of the latent phenomenon estimate, $var(\hat{\phi})$. This is an approximation of Shannon entropy that is often used to reduce computational cost. The calculated Shannon entropy exactly, $H(\hat{\phi}) - H(\hat{\phi}|x)$, is computationally expensive because it requires that $\hat{\phi}$ be estimated for each action under consideration. When using a GP belief model, this requires a GP update for every action, for every robot, at each time step. The last piece of Algorithm 1 is the learned correlation model, $h_s$, which we detail in the next section.

### C. Modeling Correlations among Observable Spatial Fields

The core hypothesis of this work is that modeling correlations among observable spatial fields can improve information gathering. By capturing the aleatoric uncertainty in the relationship, robots can sense their own observable field and use the learned model to reduce the uncertainty of their teammates' observable spatial fields.

The goal here is to learn both $h_0$, the mapping from a sensor measurement in $z^{(0)}$ to a predicted value and variance in $z^{(1)}$, as well as $h_1$, the mapping from $z^{(1)}$ to $z^{(0)}$. Note that $h_0$ is not equivalent to the inverse of $h_1$ because both $h$ map to a variance, not just a predicted value. Learning such models requires supervised learning with input values and

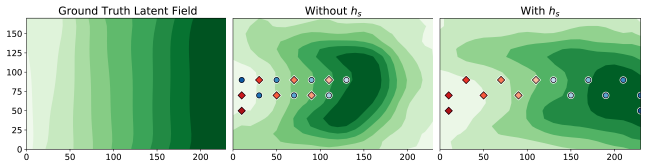



Fig. 5: *Left:* The ground truth of latent spatial field $\phi$. *Middle and right:* The belief model $\hat{\phi}$ of the latent phenomenon with robot measurements shown in red (square) and blue (circle) for the temperature and humidity robot, respectively; their sensing locations are indicated by color, with darker shades indicating later time steps. *Middle:* The correlation models $h_s$ are not used by the robot team. *Right:* When using the correlation models $h_s$, the team naturally spreads out their samples to maximize information gain.

both target values and variance. However, the data discussed in Section IV-A does not include the target variance. For example, learning the mapping from temperature to pH requires the variance of pH for a given temperature data point, but we only have *(temperature, pH)* value pairs, therefore, we transform the data. To learn $h_0$, the $d_0$ data binned and the the mean and standard deviation of the respective $d_1$ data are calculated. GP regression takes the $d_0$ bins as inputs and the $d_1$ means and standard deviations as targets. For learning $h_1$, the process is the same, flipping $d_0$ and $d_1$.

Fig. 4 shows the learned mappings $h_0$ and $h_1$ plotted over the respective historical data $d_0$ and $d_1$. These models are used in line 11. When the robot equipped with the temperature sensor takes a measurement at the point $x_t^{(0)}$, the sensor observation $z_t^{(0)}$ is used as the input to $h_0$ to update $z_t^{(1)}(x)$, the belief of pH at point $x_t^{(0)}$.

### D. Evaluation

We evaluate the performance of Algorithm 1 with and without the correlation model, i.e., we omit line 11 in the latter. As the goal of Problem III is to maximize the information with respect to the latent spatial field, the evaluation metric we use is the mean squared error between the ground truth of the latent phenomenon, $gt_2$, and the estimate $\hat{\phi}$, calculated using all points $x \in \mathcal{X}$ in the discretized space:

$$MSE = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \left( gt_2(x) - \hat{\phi}(x) \right)^2 \tag{3}$$

First, we run Algorithm 1 with the robots starting in the center of the region. The short time horizon mission in Fig. 5 shows a qualitative comparison of the algorithm's behavior with and without use of the correlation model. We see that when the correlation model is used, the robot team naturally spreads out to explore different parts of the region. Because of the correlations among the observable spatial fields, the robots are able to better estimate the expected information gain, finding that there is less information in the regions where other sensors have sampled because of the mutual information between the observable spaces.

Next, we generate 200 trials by sampling sets of random starting locations from $\mathcal{X}$. We evaluate all time horizons up to $T = 150$, reporting the MSE in Equation 3 for each time horizon. Fig. 6 shows the median MSE among the 200 trials as well as the 25th to 75th percentile range for Algorithm 1.



Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

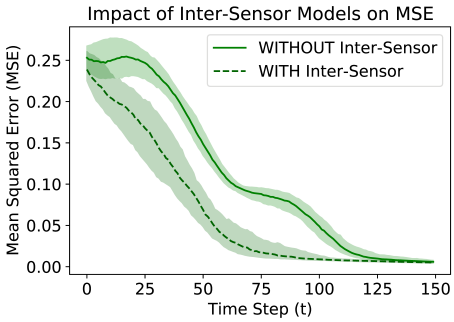



Fig. 6: The mean squared error (MSE) of the belief model $\hat{\phi}$ shows that using learned correlations among the observable spatial fields improves the performance of the adaptive sampling approach. The solid and dashed lines show the median MSE at each time step for our instantiation of Algorithm without and with correlation model $h_s$, respectively. The shaded regions show the 25th to 75th percentile range.

The comparison of this algorithm with (dotted line) and without (solid line) the learned correlations illustrates that the learned correlations among the observable spatial fields improve the efficiency of the adaptive sampling algorithm.

In the earlier phase of the sampling survey, the improvement is small because the mutual information between sensor types can not yet be exploited and the randomness of starting locations has a large impact on performance. On the other hand, during the later phase of the sampling survey, the robot samples begin to cover the entirety of the space. Therefore, both methods converge close to zero MSE and zero variance. In practice, it is prohibitively expensive to sample the entire space (and robot sampling capacity or battery life may preclude this). Therefore, adaptive sampling methods often exist in the middle ground of time horizons, where enough samples are taken to provide good information about the region, but not so many samples that it becomes overly costly. In this realm, our method performs well, using the learned correlations among the observable spatial fields to maximize the amount of information gained about the latent phenomenon in a limited number of samples.

## V. Conclusion

Real world robotic sampling methods often collect data to be analyzed offline by humans or machine learning algorithms to gain information about a latent phenomenon. However, the strength of robotic sampling is the ability for real-time adaptation. To bridge this gap, we present the ASLaP problem, where robot sampling decisions are made with respect to the information gain of a latent spatial field. We propose an informed path planning approach, ASLaP-HR to this problem using a team of heterogeneous robots.

Each robot on the team has a distinct sensor type that can measure an observable spatial field. Collectively, the measurements made by the robot team inform an estimate of the latent spatial field by using a machine learning model trained on historical data. Importantly, this learned model is used online by the robots to calculate the expected information gain of their actions with respect to the latent phenomenon. **Our method is among the firsts to explore adaptive sampling of a latent phenomenon with a team of robots using a learned model for online decision making. Additionally, our adaptive sampling strategy uses learned correlation models among heterogeneous robots to improve the performance of the team.**

We recognize that different sensor types in the real world are not independent but rather have mutual information that is a complex unknown function of their measurements. Therefore, we hypothesize that these heterogeneous robots can improve their information gathering by modeling this mutual information. We use historical sensor data to learn correlations among the observable spatial fields, and we use these learned models in our informed path planning approach. To test the impact of these correlation models, we create a simulated environment of two observable spatial fields (temp. and pH) and a latent spatial field (ORP). We run our algorithm in this simulated environment and calculate the mean squared error between the latent spatial field estimate returned by our algorithm and the ground truth latent spatial field. We aggregate our results across hundreds of randomized starting locations, showing that the learned correlations improve the efficiency of our ASLaP-HR approach.

## VI. Future Directions

As seen in Section I-A, there are myriad approaches to adaptive sampling and information gain with robot teams. The many dimensions of these approaches can be incorporated into our latent sampling framework; this includes complex motion primitives, long planning horizons, time-varying phenomena, objectives such as seeking the maximum of the latent spatial field, and uncertainty in sensors or motion, to name a few. While these open problems are shared among the adaptive sampling community, latent adaptive sampling brings new open problems. For example, new learning methods are needed for the mapping $g$ to better capture aleatoric uncertainty in the relationship between observable and latent spatial fields; current supervised learning methods focus on aleatoric uncertainty in classification, with less work on regression. An important next step that this paper enables is further investigation of the correlation models. Specifically, these systems would benefit from online learning or continual learning of these models. Such approaches would require the robots to navigate the exploration-exploitation trade-off with respect to the correlation models. In this work, we assume full communication between the robots. One near future extension of this work will be to consider varying communication ranges and intermittent communication [2].

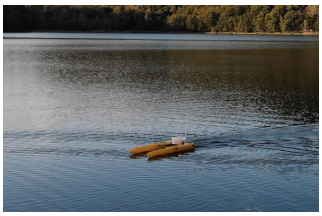
## Appendix

We make public the data set used in our second set of experiments. Data consist of water quality measurements collected at Lac Hertel in Mont Saint-Hilaire, Quebec. This data is collected using a multiparameter Sonde, Exo1 from YSI. The sensor was mounted on an autonomous surface vehicle to collect these data measurements. The data includes geotagged and timestamped measurements of 10 water quality parameters measured using four sensors: temperature, conductance, salinity, pH, oxidation reduction potential (ORP), depth, two turbidity measures and two optical

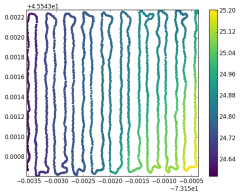


Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.

dissolved oxygen measures. We are making this data public through the following Github repository: https://github.com/SandeepManjanna/water_quality_dataset



(a) ASV collecting data at Lac Hertel



(b) Temperature measurements

Fig. 7

## 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, publisher: IEEE.

[2] S. Manjanna, M. A. Hsieh, and G. Dudek, "Scalable Multi-Robot System for Non-myopic Spatial Sampling," *arXiv preprint arXiv:2105.10018*, 2021.

[3] K. Barker and T. H. Olthof, "Relationships between nematode population densities and crop responses," *Annual Review of Phytopathology*, vol. 14, no. 1, pp. 327–353, 1976, publisher: Annual Reviews 4139 El Camino Way, PO Box 10139, Palo Alto, CA 94303-0139, USA.

[4] A. Quattrini Li, P. K. Penumarthi, J. Banfi, N. Basilico, J. M. O'Kane, I. Rekleitis, S. Nelakuditi, and F. Amigoni, "Multi-robot online sensing strategies for the construction of communication maps," *Autonomous Robots*, vol. 44, no. 3, pp. 299–319, 2020, publisher: Springer.

[5] R. Almadhoun, T. Taha, L. Seneviratne, and Y. Zweiri, "A survey on multi-robot coverage path planning for model reconstruction and mapping," *SN Applied Sciences*, vol. 1, no. 8, pp. 1–24, 2019, publisher: Springer.

[6] S. Manjanna, A. Q. Li, R. N. Smith, I. Rekleitis, and G. Dudek, "Heterogeneous Multi-Robot System for Exploration and Strategic Water Sampling," in *2018 IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 4873–4880.

[7] 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, publisher: Springer.

[8] T. Salam and M. A. Hsieh, "Heterogeneous robot teams for modeling and prediction of multiscale environmental processes," *arXiv preprint arXiv:2103.10383*, 2021.

[9] T. Manderson, S. Manjanna, and G. Dudek, "Heterogeneous robot teams for informative sampling," *arXiv preprint arXiv:1906.07208*, 2019.

[10] Y. Shi, N. Wang, J. Zheng, Y. Zhang, S. Yi, W. Luo, and K. Sycara, "Adaptive informative sampling with environment partitioning for heterogeneous multi-robot systems," in *2020 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 718–11 723.

[11] A. Chlingaryan, S. Sukkarieh, and B. Whelan, "Machine learning approaches for crop yield prediction and nitrogen status estimation in precision agriculture: A review," *Computers and electronics in agriculture*, vol. 151, pp. 61–69, 2018, publisher: Elsevier.

[12] R. Marchant, F. Ramos, S. Sanner, and others, "Sequential Bayesian optimisation for spatial-temporal monitoring." in *UAI*, 2014, pp. 553–562.

[13] R. Oliveira, L. Ott, and F. Ramos, "Bayesian optimisation under uncertain inputs," in *The 22nd Intl. conf. on artificial intelligence and statistics*. PMLR, 2019, pp. 1177–1184.

[14] Y. T. Tan, A. Kunapareddy, and M. Kobilarov, "Gaussian process adaptive sampling using the cross-entropy method for environmental sensing and monitoring," in *2018 IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6220–6227.

[15] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Proceedings of the second intl. Conf. on Autonomous agents*, 1998, pp. 47–53.

[16] I. Lluvia, E. Lazkano, and A. Ansuategi, "Active mapping and robot exploration: A survey," *Sensors*, vol. 21, no. 7, p. 2445, 2021, publisher: Multidisciplinary Digital Publishing Institute.

[17] M. Popović, T. Vidal-Calleja, J. J. Chung, J. Nieto, and R. Siegwart, "Informative path planning for active field mapping under localization uncertainty," in *2020 IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 751–10 757.

[18] M. Tzes, Y. Kantaros, and G. J. Pappas, "Distributed sampling-based planning for non-myopic active information gathering," in *2021 IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 5872–5877.

[19] Y. Kantaros, B. Schlotfeldt, N. Atanasov, and G. J. Pappas, "Sampling-based planning for non-myopic multi-robot information gathering," *Autonomous Robots*, vol. 45, no. 7, pp. 1029–1046, 2021, publisher: Springer.

[20] T. Salam and M. A. Hsieh, "Adaptive Sampling and Reduced-Order Modeling of Dynamic Processes by Robot Teams," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 477–484, 2019.

[21] N. Fung, J. Rogers, C. Nieto, H. I. Christensen, S. Kemna, and G. Sukhatme, "Coordinating multi-robot systems through environment partitioning for adaptive informative sampling," in *2019 Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3231–3237.

[22] S. Bai, T. Shan, F. Chen, L. Liu, and B. Englot, "Information-Driven Path Planning," *Current Robotics Reports*, vol. 2, no. 2, pp. 177–188, 2021, publisher: Springer.

[23] K. H. Low, J. Dolan, and P. Khosla, "Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing," in *Proceedings of the Intl. conf. on automated planning and scheduling*, vol. 19, 2009, pp. 233–240.

[24] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser, "Efficient informative sensing using multiple robots," *Journal of Artificial Intelligence Research*, vol. 34, pp. 707–755, 2009.

[25] K. H. Low, J. M. Dolan, and P. Khosla, "Adaptive multi-robot wide-area exploration and mapping," in *Proceedings of the 7th Intl. Joint Conf. on Autonomous agents and Multiagent systems-Volume 1*, 2008, pp. 23–30.

[26] S. Manjanna, H. V. Hoof, and G. Dudek, "Policy search on aggregated state space for active sampling," in *Intl. Symposium on Experimental Robotics*. Springer, 2018, pp. 211–221.

[27] P. Chand and D. A. Carnegie, "Mapping and exploration in a hierarchical heterogeneous multi-robot system using limited capability robots," *Robotics and autonomous Systems*, vol. 61, no. 6, pp. 565–579, 2013, publisher: Elsevier.

[28] X. Cai, B. Schlotfeldt, K. Khosoussi, N. Atanasov, G. J. Pappas, and J. P. How, "Non-Monotone Energy-Aware Information Gathering for Heterogeneous Robot Teams," in *2021 IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8859–8865.

[29] A. A. R. Newaz, T. Alam, J. Mondello, J. Johnson, and L. Bobadilla, "Multi-Robot Information Gathering Subject to Resource Constraints," in *2021 30th IEEE Intl. Conf. on Robot & Human Interactive Communication (RO-MAN)*. IEEE, 2021, pp. 1–6.

[30] G. Notomista, C. Pacchierotti, and P. R. Giordano, "Multi-robot persistent environmental monitoring based on constraint-driven execution of learned robot tasks," in *ICRA 2022-IEEE Intl. Conf. on Robotics and Automation*, 2022.

[31] S. Arora and S. Scherer, "Randomized algorithm for informative path planning with budget constraints," in *2017 IEEE Intl. Conf. on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 4997–5004.

[32] K. Tiwari and N. Y. Chong, *Multi-Robot Exploration for Environmental Monitoring: The Resource Constrained Perspective*. Academic Press, 2019.

[33] W. Chen and L. Liu, "Multi-objective and model-predictive tree search for spatiotemporal informative planning," in *2019 IEEE 58th Conf. on Decision and Control (CDC)*. IEEE, 2019, pp. 5716–5722.

[34] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The Intl. Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014, publisher: SAGE Publications Sage UK: London, England.

[35] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Machine Learning*, vol. 110, no. 3, pp. 457–506, 2021, publisher: Springer.

[36] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies." *Journal of Machine Learning Research*, vol. 9, no. 2, 2008.




Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 14:45:19 UTC from IEEE Xplore. Restrictions apply.