

Adaptive sampling of a stationary Gaussian spatial process by a team of robots with heterogeneous dynamics and measurement noise variance

Michael Brancato¹ and Artur Wolek¹

¹Department of Mechanical Engineering and Engineering Science, The University of North Carolina at Charlotte

April 18, 2024

Adaptive sampling of a stationary Gaussian spatial process by a team of robots with heterogeneous dynamics and measurement noise variance

Michael Brancato and Artur Wolek
Department of Mechanical Engineering and Engineering Science
The University of North Carolina at Charlotte

Abstract

This paper presents an algorithm for a team of heterogeneous mobile robots to estimate and adaptively sample a stationary, isotropic, Gaussian process. An estimation framework is proposed to assimilate measurements from robots with differing sensing capabilities (i.e., measurement noise variance). To improve computational efficiency of the Gaussian process regression, the survey area is divided into regions that each use a common semivariogram matrix constructed with a truncated measurement set determined by an adaptive selector. As the mission proceeds, a Voronoi-based algorithm periodically partitions a density function—representing time-varying sampling priority—to identify high-value sampling locations. The path for each robot is modeled as a mechanical system: a sequence of masses (waypoints) are interconnected by springs and dampers and pulled towards Voronoi cell centroids. At each path planning cycle the robots are iteratively simulated with heterogeneous dynamics (e.g., speed, turn radius) following their respective waypoint paths as stiffness/damping parameters are adjusted to satisfy mission time constraints. Numerical simulations show that the proposed approach reduces mapping error when compared to non-adaptive lawnmower coverage of a survey area. The algorithm is demonstrated experimentally using two cooperating autonomous surface vessels to map the bathymetry in a section of Lake Norman near Charlotte, NC.

1 Introduction

Teams of robots equipped with sensors can coordinate their actions to cover a large area and maximize information gain using adaptive sampling (AS) algo-

rithms [1], also known as informative path planning (IPP) [2]. Robotic sensor networks provide an automated, efficient, and scalable means to collect continuous spatial data in applications such as precision agriculture, terrain mapping, and environmental monitoring [3–7]. Adaptive sampling algorithms use a model of the spatiotemporal process to assimilate measurements and predict the underlying attribute of interest in unsampled regions. Often the spatiotemporal distribution of the attribute can be modeled as a Gaussian Process (GP). A GP is a random process characterized by a mean and a covariance function. GP regression (known as kriging in the geosciences [8]) enables data-driven estimation of a spatial attribute in unsampled locations as a weighted linear combination of existing measurements. Adaptive sampling algorithms that utilize GP-based models of continuous spatial distributions have been widely studied, including for robotic mapping of light intensity [7], water body salinity or conductivity [4, 9], temperature [5], terrain height [10], chemical concentrations [11], and soil moisture [3]. The main idea in AS/IPP is to allocate sampling trajectories to robots by periodically optimizing a performance metric related to the quality of the estimated spatial map of the attribute. In this way, robots can more efficiently estimate a spatially distributed quantity of interest, especially when there is not enough time to perform a more exhaustive full coverage survey.

Our aim in this paper is to design and demonstrate an adaptive sampling algorithm for heterogeneous robot teams that have continuous dynamics described by ordinary differential equations (ODEs) and to allocate sampling trajectories that are cognizant of each robot’s differing dynamics and sensing capabilities. In our method, we consider the precise path the vehicle will follow and corresponding mission time as a constraint. The agents may have differing dynamics and measure the same attribute but with varying

sensor quality (i.e., each robot’s sensor has differing measurement noise variance). A Gaussian process regression is proposed to efficiently handle this heterogeneous sensing scenario for a stationary, isotropic spatial field. Then, an adaptive sampling algorithm is developed that iteratively simulates the motion of each robot (with its heterogeneous dynamics) through a sequence of path waypoints from the current time to the final mission time to periodically replan the path. Samples are obtained at a fixed time intervals along the path and the sampling path is designed to minimize mapping error.

1.1 Contributions

The contributions of this paper are (1) an efficient Gaussian process regression framework to fuse measurements from multiple robots with heterogeneous measurement noise using an adaptive truncation with common data neighborhoods, (2) an adaptive sampling approach that considers heterogeneous vehicle dynamics and sensing to allocate sampling trajectories. The approach is compared to non-adaptive lawnmower coverage of a survey area through numerical Monte Carlo simulations. The feasibility of the method is demonstrated by a field experiment in which two autonomous surface vessels adaptively map the bathymetry in a subsection of a lake.

1.2 Organization

The paper is organized as follows. Section 1.3 presents related work. Section 2 introduces notation and mathematical background on Gaussian process regression and centroidal Voronoi tessellations. Section 3 formulates the optimization problem by describing the robot dynamics, sensing, and the cost function and constraints. Section 4 proposes an efficient approach for Gaussian process regression with heterogeneous measurements. Section 5 describes the adaptive sampling algorithm. Section 6 provides results from simulations. Section 7 discusses a real-world demonstration of the algorithm. The paper is concluded in Section 8.

1.3 Related Work

The following section highlights prior work related to AS/IPP algorithms and compares and contrasts how heterogeneity, vehicle dynamics, and spatial process estimation has been incorporated into prior problem formulations. Various techniques used for optimizing the sampling path and the relation of the proposed approach to prior work is discussed.

Heterogeneity Many prior works have focused on AS/IPP for a single vehicle [3–5, 7, 10, 12–14] or a set of homogeneous vehicles [9, 15–27]. When increasingly larger teams of robots are considered communication constraints [15, 28–30] and scalability [21, 31] must be addressed. Heterogeneity arise when multiple vehicles have differing platform dynamics and/or sensing capabilities. In [28], vehicles with sensors that have varying fields of view are tasked with inferring the state of an environment. In [32], agents use differing sensor types to estimate a latent phenomena that is correlated with the unique sensor measurements. In [33] heterogeneous vehicles are equipped with unique sensor suites and use a three-phase process to scan a region and identify points of interest. In [31] a general framework to consider multiple heterogeneity criteria (speed, battery life, terrain traversability, sensor type etc.) is developed using a heterogeneity cost space to leverage robot capabilities. Other examples of heterogeneity in the context of GP-based mapping include AS/IPP while performing another task [34], for front detection [35], and for symbiotic planning along a route with a second supporting platform to conserve battery life [36]. AS/IPP and mapping/exploration methods that consider heterogeneity outside the context of GP models of spatial fields include [28, 37–40].

Models of Vehicle Dynamics To reduce complexity, a common AS/IPP assumption is that the vehicle moves along straight segments at a constant speed without any turning dynamics [3–5, 7, 9, 12–14, 17, 19–21, 25–27, 32]. Motion may be in arbitrary directions, constrained to certain directions (e.g., in the four cardinal directions) [21], constrained between certain waypoints [32], or constrained to follow continuous paths that are produced by the optimizer [15]. Ignoring vehicle turn dynamics is well justified when the spatial length scale of the survey is much larger than the length scale of vehicle turning maneuvers. Vehicle low-level controllers are assumed to adapt to any environment factors (e.g., wind, current) and account for other dynamic model complexities. However, when operating in more constrained environments, or when accurate accounting of energy expenditure or time is critical, a more sophisticated dynamic model may need to be incorporated into the planning process. A few authors have developed algorithms that explicitly utilize an ordinary differential equation (ODE) model of the robot’s motion. In [10] the motion of an autonomous surface vessel is parametrized using a kinematic model with speed and turn-rate constraints. In [28] continuous time integrator dynamics are used to approximate the motion of a quadrotor with a position controller

Ref.	Multi-robot	GP	Heterog. Sensing	Heterog. Motion	Robot ODE	Solution Type
Manjana et al. [21]	x	x	–	–	–	waypoint/discrete sequence
Malencia et al. [32]	x	x	x (sensor type)	–	–	waypoint/discrete sequence
Kathen et al. [33]	x	x	x (sensor type)	–	–	waypoint/discrete sequence
Salam et al. [40]	x	–	x (sensor fidelity)	x	–	waypoint/discrete sequence
Shi et al. [31]	x	x	x (generalizable)	x	–	coverage configuration
Zhang et al. [24]	x	x	–	–	x	continuous path
Tan et al. [10]	–	x	–	–	x	continuous path
Julian et al. [28]	x	–	x (sensor FOV)	–	x	continuous path
This paper	x	x	x (sensor noise)	x	x	continuous path

Table 1: Comparison of this paper to related work that considers heterogeneity and/or the ordinary differential equation (ODE) describing the robot’s dynamics.

following waypoints. In [24], three-dimensional non-holonomic rigid body underwater vehicle dynamics are considered. Other works that, in principle, could incorporate vehicle dynamics include [22,23] using B-splines, or [14,41] using the steering subroutine of a sampling-based motion planning algorithm.

Estimation approaches GP regression and hyperparameter optimization is widely used to estimate a spatiotemporal process of interest from noisy measurements [5, 7, 9, 10, 13, 16, 17, 19, 20] [23–25]. GP regression has been adapted to consider special cases, such as localization uncertainty [11], distributed estimation [42], correlation among several distinct measured quantities and a latent process [32], mixtures of GP models [26, 31], and when the variance depends on the measurement itself [3]. However, one drawback of GP regression (especially for robotic systems with limited onboard computation) is that the time-complexity is $O(N^3)$ [43, ch.8] where N is the number of samples. To address this challenge GP-based adaptive sampling algorithms with an iterated covariance update [13], truncated measurements [15], and fixed-complexity [44] have been proposed. Other approaches for improving scalability use global or local approximations such as sparse GPs [24, 45]. Sparse GPs reduce the computational burden and can account for time-varying environments by periodically re-estimating the hyperparameters [4] and can be used for streaming data [22] and distributed data [18].

Non-GP-based approaches have also been adopted to model spatiotemporal processes for adaptive sampling, for example, using Kalman filters to estimate the weights of a set of basis functions [16, 41] or using sequential Bayesian filters [28]. In the absence of a spatiotemporal process model and estimation method, exhaustive or partial coverage paths techniques that consider energy limitations of the robotic team [12, 46, 47] can be used for information gathering.

Optimization Problem Formulation and Approach

Prior work has considered numerous optimization objectives and techniques. When the planning space is sufficiently small each candidate sampling location can be evaluated to identify the optimal action at each iteration [3, 5, 7, 32, 32]. In [5] an objective function balances reduction in variance and gradient intensity and searches a waypoint graph to determine the next best survey line. In [7], Bayesian optimization selects the best sampling location from a set using a metric that balances the mean and variance of prediction with distance to travel. In [3] both greedy and weighted randomized (Monte Carlo) selection criteria are used to determine sampling points with high kriging variance.

The adaptive sampling problem can also be formulated as a combinatorial-optimization to visit a particular set of points in an optimal sequence. In [36], prior knowledge is used to identify potentially mislabeled (PML) points and a variant of a traveling salesperson problem (TSP) is formulated to visit either all of the points or maximizing the number of PML points visited subject to energy constraints. In [4] the sampling locations are chosen from a pre-defined grid to maximize mutual information using the predicted mean and variance at each sample location. In [22] the inducing points of a sparse GP are used as waypoints to formulate an assignment problem.

Gradient-following techniques have also been investigated to maximize mutual information [28], balance energy expenditure with minimizing error variances at target positions [16], to minimize the average of the prediction error variances at target positions [15], or to minimize the posterior entropy after making a new observation [13]. Greedy approximations that leverage spatial decompositions have also been proposed to optimize mutual information [9] or to optimize condition entropy [17]. Receding horizon cross-entropy trajectory optimization has investigated to maximize an upper confidence bound [10]. Techniques that use coverage control and geometric optimization (e.g.,

Voronoi partitions) have also been applied [17–19, 26]. Other methods include using Markov Decision Process (MDP) [21] with a policy gradient algorithm to maximize the accumulated reward derived from a spatial adaptive scoremap, reinforcement learning to reduce uncertainty balanced with trajectory cost [20], evolutionary algorithms [23, 24], and sampling-based methods [14, 41].

Relation of Proposed Approach to Prior Work

Table 1 highlights several prior works that have investigated AS/IPP algorithms that consider (1) heterogeneous vehicle motion and sensing, and (2) differential constraints governing vehicle motion, and (3) time or energy constraints computed by considering the full continuous path motion of the robotic system. However, to our knowledge, no prior work considers all three aspects simultaneously. This paper contributes to the literature and methodologies available to tackle such problems. Specifically, our paper builds on the work in [12], which is an offline planning method for a single-vehicle to maximize coverage with energy-constrained robots. Our approach extends [12] to a multi-robot framework that considers vehicle dynamics. We adopt the mass-spring-damper architecture used for waypoint position tuning in [12] and extend it into an adaptive framework that involves GP estimation and periodic adaptation and path re-planning. Rather than using a standard Voronoi diagram as in [12], our method uses a centroidal Voronoi partition with a time-varying surface (that changes as measurements are assimilated) to represent sampling priority and identify high-value sampling locations that attract path waypoints. Moreover, the process of tuning spring and damper constants in [12] is extended to allow for satisfying a mission time constraint computed for the continuous dynamics of each robot. To our knowledge this paper is the first to make a connection between the two geostatistical techniques of Common Data Neighborhoods (CDN) [48] and heterogeneous measurement-error filtered kriging (HFK) [49] with robotic informative path planning.

2 Background

This section introduces notation used throughout the paper and relevant background, including: Gaussian spatial processes, ordinary and heterogeneous filtered-error kriging, and centroidal Voronoi tessellations.

2.1 Gaussian Spatial Processes

A Gaussian process (GP) is a infinite collection of joint random variables [50] from which any finite subset have a joint Gaussian distribution. A GP is denoted by [43]

$$Z(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) , \quad (1)$$

where the collection of random variables $Z(\mathbf{x})$ have a mean function $\mu(\mathbf{x}) = E[Z(\mathbf{x})]$ and covariance function (or kernel) $k(\mathbf{x}, \mathbf{x}') = E[(Z(\mathbf{x}) - \mu(\mathbf{x}))(Z(\mathbf{x}') - \mu(\mathbf{x}'))]$, $E[\cdot]$ is the expected value of a random variable, and $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$ are two locations in the input space (e.g., a spatial domain). Gaussian processes that are second-order stationary have a constant mean $\mu(\mathbf{x}) = m$ and a covariance function between points $\mathbf{x}$ and $\mathbf{x}' = \mathbf{x} + \mathbf{h}$ that is invariant under translation [8, 49]

$$k(\mathbf{x}, \mathbf{x} + \mathbf{h}) = E[(Z(\mathbf{x}) - m)(Z(\mathbf{x} + \mathbf{h}) - m)] = k(\mathbf{h}) , \quad (2)$$

where $\mathbf{h}$ is known as the lag between two points $\mathbf{x}$ and $\mathbf{x}'$. Covariance functions describe the amount of variation of the process and the degree to which two points in the spatial process are related (i.e., the smoothness of the process). This paper uses the isotropic Gaussian covariance function

$$k(\mathbf{h}) = \sigma_0^2 e^{-3(\|\mathbf{h}\| \omega^{-1})^2} \quad (3)$$

where σ_0^2 is the overall variance of the process and ω is the length-scale of the GP. The factor of three in the exponent of (3) ensures that covariance function is $0.05\sigma_0^2$ when $\|\mathbf{h}\| = \omega$, a standard feature in kriging for geosciences [8, p.78].

When sampling a real-world scalar field, the covariance function of the model is generally unknown and must be estimated. Often the covariance function is substituted with the semivariogram, γ , which is a measure of dissimilarity of measurements with spatial distance [8]. For GPs which are second-order stationary, the covariance and semivariogram are related by [8]

$$\gamma(\mathbf{h}) = k(\mathbf{0}) - k(\mathbf{h}) . \quad (4)$$

The semivariogram can be estimated by fitting a curve to a scatter-plot of the average variation in measurements over a given distance [8]. The fit curve is termed the experimental semivariogram, $\hat{\gamma}(\mathbf{h})$.

2.2 Ordinary Kriging

Gaussian processes can be conditioned on prior measurements to make predictions at unsampled points of interest. A type of GP regression developed for use

in geostatistics is known as ordinary kriging [8]. Ordinary kriging is referred to as a Best, Linear Unbiased Estimator (BLUE) because it minimizes mean square error using a weighted combination of measurements for each estimate. Let $\mathbf{Z}(\mathbf{X}) = [z_1, \dots, z_M]^T$ be a vector of noise-free measurements of the random function $Z(\mathbf{x})$ at spatial locations $\mathbf{X} = \{\mathbf{x}_i \mid i = \{1, 2, \dots, M\}\}$. The optimal weights in ordinary kriging, $\boldsymbol{\lambda}_{\text{OK}}$, and the Lagrange multiplier, μ_{OK} , for a desired a prediction point, $\mathbf{x}_0$, are [49]

$$\begin{bmatrix} \boldsymbol{\lambda}_{\text{OK}} \\ \mu_{\text{OK}} \end{bmatrix} = \begin{bmatrix} \hat{\Gamma}(\mathbf{X})\mathbf{Z} & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} \hat{\gamma}(\mathbf{X}, \mathbf{x}_0) \\ 1 \end{bmatrix}, \quad (5)$$

where the elements of the semivariogram vector $\hat{\gamma}(\mathbf{X}, \mathbf{x}_0) \in \mathbb{R}^{M \times 1}$ are $\hat{\gamma}(\mathbf{X}, \mathbf{x}_0)_i = \hat{\gamma}(\|\mathbf{x}_i - \mathbf{x}_0\|)$, $\mathbf{1} \in \mathbb{R}^{M \times 1}$ is a column of ones, and the semivariogram matrix $\hat{\Gamma}(\mathbf{X}) \in \mathbb{R}^{M \times M}$ has elements [49]

$$\hat{\Gamma}(\mathbf{X})_{ij} = \hat{\gamma}(\|\mathbf{x}_i - \mathbf{x}_j\|) \quad (6)$$

for $\mathbf{x}_i, \mathbf{x}_j \in \mathbf{X}$. The estimate and variance of the estimation error at the prediction point $\mathbf{x}_0$ are [49]

$$\hat{Z}_{\text{OK}}(\mathbf{x}_0) = \boldsymbol{\lambda}_{\text{OK}}^T \mathbf{Z}(\mathbf{X}) \quad (7)$$

$$\sigma_{\text{OK}}^2(\mathbf{x}_0) = \boldsymbol{\lambda}_{\text{OK}}^T \hat{\gamma}(\mathbf{X}, \mathbf{x}_0) + \mu_{\text{OK}}, \quad (8)$$

respectively. This estimation can be repeated for various spatial locations to produce an estimate of the underlying GP. Figure 1 shows the result of kriging over a set of uniformly spaced grid of prediction points.

2.3 Heterogeneous Error Filtered Kriging

Ordinary kriging is an exact interpolator in the sense that the prediction at sampled points is exactly equal to the measurement observed at that sample point [49]. However, when the measurements are corrupted by noise the filtered Kriging approach is used to produce a smoothed prediction of the measurement-error-free value of the process [49]. Heterogeneous measurement-error filtered kriging (HFK) extends ordinary kriging to the case of site-specific measurement error with known variance [49]. Let

$$\tilde{Z}(\mathbf{x}_i) = Z(\mathbf{x}_i) + \epsilon(\mathbf{x}_i), \quad (9)$$

represent the noise-corrupted measurement of the attribute Z at location $\mathbf{x}_i$ where $\epsilon(\mathbf{x}_i)$ is the realization of a zero-mean Gaussian random variable with site-specific measurement noise variance $\sigma_{\eta}^2(\mathbf{x}_i)$. The vector of noise-corrupted measurements obtained at locations $\mathbf{X}$ according to (9) is $\tilde{\mathbf{Z}}(\mathbf{X})$. This site-specificity




Figure 1: Two-dimensional field estimation using ordinary kriging. The surface representing $Z(\mathbf{x})$ is characterized by a covariance function (3) with $\boldsymbol{\theta} = [\zeta, \omega, \sigma_0^2]^T = [0, 0.3, 1]^T$ and with a mean $\mu = 5$. A total of 20 measurements, represented by the red projections from the surface to the x_1 - x_2 -plane, are the measurements. The coloring on the x_1 - x_2 -plane indicates estimation variance; the darker the color, the lower the relative variance. At locations where measurements were taken, the variance equals zero since the measurements are noise-free.

will be leveraged to account for heterogeneity among robots in Sec. 4. Christensen [49] proposed to account for noisy measurements (9) using the HFK estimator

$$\hat{Z}_{\text{HFK}}(\mathbf{x}_0) = \boldsymbol{\lambda}_{\text{HFK}}^T \tilde{\mathbf{Z}}(\mathbf{X}) \quad (10)$$

$$\hat{\sigma}_{\text{HFK}}^2(\mathbf{x}_0) = \tilde{\gamma}^T \boldsymbol{\lambda}_{\text{HFK}}(\mathbf{X}, \mathbf{x}_0) + \mu_{\text{HFK}}, \quad (11)$$

where the optimal kriging weights, $\boldsymbol{\lambda}_{\text{HFK}}$, and Lagrange multiplier, μ_{HFK} , are found from

$$\begin{bmatrix} \boldsymbol{\lambda}_{\text{HFK}} \\ \mu_{\text{HFK}} \end{bmatrix} = \begin{bmatrix} \tilde{\Gamma}(\mathbf{X}) & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} \tilde{\gamma}(\mathbf{X}, \mathbf{x}_0) \\ 1 \end{bmatrix}. \quad (12)$$

In (12), $\tilde{\Gamma}(\mathbf{X})$ and $\tilde{\gamma}(\mathbf{X}, \mathbf{x}_0)$ are augmented semivariogram matrices and vectors, respectively. The matrix $\tilde{\Gamma}(\mathbf{X})$ is defined by augmenting (6) with the average site-specific variance along off-diagonal elements [49]

$$\tilde{\Gamma}(\mathbf{X})_{ij} = \hat{\gamma}^*(\|\mathbf{x}_i - \mathbf{x}_j\|) + (1 - \delta_{\|\mathbf{h}_{ij}\|}) \frac{\sigma_{\eta}^2(\mathbf{x}_i) + \sigma_{\eta}^2(\mathbf{x}_j)}{2}, \quad (13)$$

where $\delta_{\|\mathbf{h}\|} = 1$ when $\|\mathbf{h}\| = 0$, and $\delta_{\|\mathbf{h}\|} = 0$ otherwise, and $\tilde{\gamma}(\mathbf{X}, \mathbf{x}_0)$ is defined by augmenting the i th element of the semivariogram vector with half the site-specific variance of the measurement [49]

$$\tilde{\gamma}(\mathbf{X}, \mathbf{x}_0)_i = \hat{\gamma}^*(\|\mathbf{x}_i - \mathbf{x}_0\|) + \frac{\sigma_{\eta}^2(\mathbf{x}_i)}{2}. \quad (14)$$

In (13) and (14), $\hat{\gamma}^*$ refers to the estimated semivariogram of the noise-free process $\hat{\gamma}^*$. In practice, the

noisy observations $\tilde{Z}$ are used to estimate the semivariogram $\hat{\gamma}$ of the noise-corrupted process and a *nugget* term ζ is used to account for a discontinuity in the covariance function at $\mathbf{h} = 0$ to model micro-scale variation/sensor measurement error [50]. The semivariogram model used here follows from (3) and (4) with the addition of a nugget term

$$\hat{\gamma}(\|\mathbf{h}\|; \boldsymbol{\theta}) = \zeta(1 - \delta_{\|\mathbf{h}\|}) + \sigma_0^2 \left(1 - e^{-3(\|\mathbf{h}\| \omega^{-1})^2}\right), \quad (15)$$

and the vector of hyper-parameters is $\boldsymbol{\theta} = [\zeta, \omega, \sigma_0^2]^T$. In [49] it is proposed that the noise-free estimated semivariogram $\hat{\gamma}^*$ can be obtained from the noise-corrupted estimated semivariogram $\hat{\gamma}$ by subtracting the average of the measurement-error variance across all M sampling locations from the nugget estimated using the noisy measurements [49], i.e., $\hat{\gamma}^*(\|\mathbf{x}_i - \mathbf{x}_j\|) = \hat{\gamma}(\|\mathbf{x}_i - \mathbf{x}_j\|; \zeta^*, \omega, \sigma_0^2)$ where

$$\zeta^* = -\frac{1}{M} \sum_{k=1}^M \sigma_\eta^2(\mathbf{x}_k). \quad (16)$$

2.4 Centroidal Voronoi Tessellations

This work uses centroidal Voronoi tessellations (CVTs) to identify high-value sampling locations. Recall that a standard Voronoi tessellation (VT) partitions a planar region, $Q \subset \mathbb{R}^2$, into a set containing k non-intersecting polygonal cells, $\{V_i\}_{i=1}^k$, whose union equals Q . Given a set of generating points $\{\mathbf{g}_i\}_{i=1}^k \in Q^k$, the Voronoi cell V_i is the set of points within Q that are closer to the cell's generating point, $\mathbf{g}_i$, than any other generating point [51]:

$$V_i = \{\mathbf{x} \in Q \mid \|\mathbf{x} - \mathbf{g}_i\| \leq \|\mathbf{x} - \mathbf{g}_j\|\} \quad (17)$$

for all $j = \{1, 2, \dots, k\}$ with $j \neq i$. When a density function is defined over the domain $\rho : Q \rightarrow \mathbb{R}$, each cell has a mass and center of mass

$$M_{V_i} = \int_{\mathbf{x} \in V_i} \rho(\mathbf{x}) d\mathbf{x} \quad \text{and} \quad \mathbf{c}_{V_i} = \frac{1}{M_{V_i}} \int_{\mathbf{x} \in V_i} \mathbf{x} \rho(\mathbf{x}) d\mathbf{x}, \quad (18)$$

respectively. CVTs [51] are variants of VTs in which the generating points of the VT are co-located with the associated cell's centroid (i.e., $\mathbf{g}_i = \mathbf{c}_{V_i}$). In this work, (18) is discretized by replacing the integrals with summations over a uniform grid, and the CVT is computed using a variant of Lloyd's method [52].

3 Problem Formulation

This section describes the robot dynamics and sensing and formulates the cost function and optimization problem of the paper.

3.1 Heterogeneous Agent Dynamics

Consider a set of N robots with heterogeneous dynamics. Each robot is modeled as a differential thrust autonomous surface vessels with equations of motion

$$\begin{aligned} \ddot{x}_i &= \frac{1}{m_i} \{[u_r + u_l] \cos(\theta_i) - b_i \dot{x}_i\} \\ \ddot{y}_i &= \frac{1}{m_i} \{[u_r + u_l] \sin(\theta_i) - b_i \dot{y}_i\} \\ \ddot{\theta}_i &= \frac{L_i}{2I_i} (u_r - u_l), \end{aligned} \quad (19)$$

where $(x_i, y_i) \in \mathbb{R}^2$ is the planar position relative to a fixed inertial reference frame with origin O and orthonormal basic vectors $\{\mathbf{i}_1, \mathbf{i}_2, \mathbf{i}_3\}$ and $\theta_i \in [0, 2\pi)$ is the heading (see Fig. 2) for the i th robot with $i \in \{1, 2, \dots, N\}$. This choice of dynamics is motivated by our experimental work (see Sec. 7) but it can be generalized to other vehicle models. In (19), m_i is the mass, I_i is the rotational inertia, L_i is the distance between thrusters, and b_i is a damping coefficient (i.e., linear drag) of the i th robot. The left and right thruster forces, u_l and u_r , respectively, are bounded: $u_l, u_r \in [u_{\min}, u_{\max}]$ where $u_{\min}$ and $u_{\max}$ are a minimum and maximum thrust force. The i th robot follows a reference path consisting of a set of waypoints $\mathbf{W}_i = \{\mathbf{w}_{i,1}, \mathbf{w}_{i,2}, \dots, \mathbf{w}_{i,N_w}\}$ where $\mathbf{w}_{i,j} \in Q$ for all $j = \{1, 2, \dots, N_w\}$ and N_w is the number of waypoints in the path. For a given waypoint, $\mathbf{w}_{i,j} = (x_w, y_w) \in \mathbf{W}_i$, the homing guidance law $\theta_d = \text{atan2}(y_w - y_i, x_w - x_i)$ gives a desired heading and the thrusters are commanded according to

$$\begin{aligned} u_l &= \text{sat}(\delta_v - \delta_\theta; 0, u_{\max}) \\ u_r &= \text{sat}(\delta_v + \delta_\theta; 0, u_{\max}), \end{aligned} \quad (20)$$

where δ_v and δ_θ are the outputs of PID controllers tracking a desired speed v_d and the desired heading θ_d , respectively, and $\text{sat}(\kappa; a_l, a_h)$ is the saturation function that bounds an argument κ between a lower bound a_l and upper bound a_h such that $a_l \leq \kappa \leq a_h$. Once the position of the i th robot is sufficiently close to the j th waypoint, $\|[\mathbf{x}_i, \mathbf{y}_i]^T - \mathbf{w}_{i,j}\| \leq R_c$, where $R_c \in \mathbb{R}$ is a capture radius, the next waypoint in the sequence becomes active by incrementing j . The closed-loop dynamics of the i th robot can be summarized as

$$\dot{\mathbf{s}}_i = \mathbf{f}_i(\mathbf{s}_i, \mathbf{u}_i(\mathbf{s}_i); \mathbf{W}_i), \quad (21)$$

where $\mathbf{s}_i = [x_i, y_i, \theta_i, \dot{x}_i, \dot{y}_i, \dot{\theta}_i]^T \in \mathbb{R}^2 \times [0, 2\pi) \times \mathbb{R}^3$ is the state of the system (19) rewritten in first-order form and $\mathbf{u}_i(\mathbf{s}_i)$ is the guidance and control law (20).




Figure 2: The i th robot is modeled as a differential thrust vehicle. The desired heading angle θ_d points the vehicle towards waypoint $\mathbf{w}_j$. Once robot's position (x, y) is within a capture radius R_c of the waypoint the next waypoint in the sequence becomes active.

3.2 Heterogeneous Agent Sensing

The team of N robots can have differing measurement noise variance (e.g., due to differing sensors or signal processing instrumentation). The measurement noise variance for each robot $\sigma_{i,\eta}^2 = E[(Z - \tilde{Z})^2]$ characterizes the expected difference between the true and measured values of the observed scalar field according to (9). The robots are assumed to be in constant communication with a centralized base station.

3.3 Cost Function

Given a waypoint list $\mathbf{W}_i$, initial conditions $\mathbf{s}_i(t_0)$, a sensing time interval T_s , the closed-loop dynamics (21), and a maximum mission time T_m , the map

$$\phi_i(\mathbf{W}_i; \mathbf{s}_i(t_0), t_0, T_m, T_s) = [\mathbf{q}_{i,1}, \mathbf{q}_{i,2}, \dots, \mathbf{q}_{i,M}] \quad (22)$$

returns the i th robot's sampling locations arranged as columns of a matrix in chronological order where $\mathbf{q}_{i,j} \in Q$ for all $j = \{1, 2, \dots, M\}$. During optimization, these sampling locations ϕ_i are obtained by numerically simulating the robotic team. Let the matrix $\mathbf{X} \in \mathbb{R}^{2N \times M}$ denote the stacked set of sampling locations across all robots at the end of the mission

$$\mathbf{X}(\mathbf{W}) = [\phi_1(\mathbf{W}_1)^T, \phi_2(\mathbf{W}_2)^T, \dots, \phi_N(\mathbf{W}_N)^T]^T. \quad (23)$$

The corresponding measurements made at each location in $\mathbf{X}$ are denoted $\tilde{\mathbf{Z}} \in \mathbb{R}^{N \times M}$.

The goal of the adaptive sampling algorithm is to design waypoint paths for each robot to collectively minimize the mapping error (ME)

$$\text{ME}(\mathbf{X}, \tilde{\mathbf{Z}}) = \frac{1}{|\mathcal{P}|} \sum_{i=1}^{|\mathcal{P}|} |\hat{Z}(\mathbf{p}_i) - Z(\mathbf{p}_i)|, \quad (24)$$

which is the average difference between the estimated field $\hat{Z}$ and the true field Z across $|\mathcal{P}|$ grid points

at which the spatial process is estimated. We consider $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_{m \cdot n}\}$ to be a uniform grid with m columns and n rows where $\mathbf{p}_i \in Q$ for all $i = \{1, 2, \dots, m \cdot n\}$.

3.4 Problem Statement

The optimization problem is to find the set of waypoints $\mathbf{W}^*$ that minimize the mapping error:

$$\arg \min_{\mathbf{W}} \text{ME}(\mathbf{X}(\mathbf{W}), \tilde{\mathbf{Z}}) \quad (25a)$$

$$\text{subject to } \dot{\mathbf{s}}_i = \mathbf{f}_i(\mathbf{s}_i, \mathbf{u}_i(\mathbf{s}_i); \mathbf{W}_i), \quad (25b)$$

$$\mathbf{s}_i(t_0) = (\mathbf{s}_i)_0, \quad (25c)$$

$$\mathbf{u}_{\min} \leq \mathbf{u}_i(\mathbf{s}_i) \leq \mathbf{u}_{\max}, \quad (25d)$$

$$T(\mathbf{w}_{i,N_w}) \leq T_m \text{ for all } i = 1, \dots, N, \quad (25e)$$

$$\tilde{Z}(\mathbf{x}_j) = Z(\mathbf{x}_j) + \epsilon(\mathbf{x}_j), \quad (25f)$$

for all $\mathbf{x}_j \in \mathbf{X}(\mathbf{W})$.

The cost (25) depends on the sampling locations $\mathbf{X}$ which are, in turn, determined by the waypoints $\mathbf{W}$ to be optimized. Heterogeneity in dynamics is encapsulated by the equations of motion (25b), the initial conditions (25c), and the minimum/maximum thrust input constraints (25d). The mission time constraint is encoded by (25e) so that the estimated capture time of the final mission waypoint is less than or equal to the desired total mission time. In (25e) the function $T(\mathbf{w}_{i,j}) = \psi(\mathbf{w}_{i,j}; \mathbf{W}_i, \mathbf{s}_i(t_0), t_0)$ denotes the time when robot i reaches its j th waypoint given its list of waypoints and the initial state. The variation between robots' sensing capability is described by (25f) which implicitly assumes robot-dependent measurement noise variance, as described in Sec. 3.2.

The above optimization problem is addressed in two parts. First, a method for Gaussian process regression is introduced in Sec. 4 that uses the measurements (25f) to allow predicting the GP over the grid of points $\mathcal{P}$. Next, Sec. 5 describes the adaptive sampling algorithm to find waypoint paths that aim to produce paths with low mapping error (25). Path optimization is performed by a centralized base station.

4 Efficient Gaussian Process Regression

In this section, a modified kriging approach is proposed to account for robots with heterogeneous measurement noise variance and to improve computational efficiency.

4.1 Filtered Kriging with Heterogeneous Sensors

Consider a robotic sensor network of N robots that measure a spatial attribute Z with heterogeneous measurement variance $\sigma_{\eta,i}^2$ for $i = 1, \dots, N$. This work proposes to specialize heterogeneous filtered kriging (HFK, see Section 2.3) to the case of heterogeneous sensing robots by replacing the site-specific measurement variance with a robot-specific one. Suppose that each robot collects measurements at the same sampling rate and begins to sample at the same time. The total number of measurements from each robot is equal and all measurements taken from all robots are used to populate the set $\tilde{\mathbf{Z}}$. An indicator function, $\beta(\mathbf{x})$, is defined that returns the measurement variance σ_{η}^2 associated with the robot that took the measurement at $\mathbf{x}$, assuming no two samples are co-located. Thus, (13) and (14) become

$$\tilde{\Gamma}(\mathbf{X})_{ij} = \hat{\gamma}^*(\|\mathbf{x}_i - \mathbf{x}_j\|) + (1 - \delta_{\|\mathbf{h}_{ij}\|}) \frac{\beta(\mathbf{x}_i) + \beta(\mathbf{x}_j)}{2} \quad (26)$$

$$\tilde{\gamma}(\mathbf{X}, \mathbf{x}_0)_i = \hat{\gamma}^*(\|\mathbf{x}_i - \mathbf{x}_0\|) + \frac{\beta(\mathbf{x}_i)}{2}, \quad (27)$$

while (16) becomes

$$\hat{\gamma}_Z(\mathbf{x}_i, \mathbf{x}_j) = \hat{\gamma}_{\tilde{\mathbf{Z}}} \left(\mathbf{x}_i, \mathbf{x}_j; \zeta_{\tilde{\mathbf{Z}}} - \frac{1}{N} \sum_{i=1}^N \sigma_{\eta,i}^2, \omega, \sigma_0^2 \right). \quad (28)$$

4.2 Adaptive Spatial Truncation

The time complexity of kriging is dominated by the inversion of the semivariogram matrix in (5) or (12). To improve efficiency, spatial truncation can be used to reduce the size of this matrix by considering only nearby measurements with strong correlations to an estimation point and rejecting measurements that have little influence [8], [15]. Typical truncation methods rely on geometric selector regions (squares, circles, etc.) centered about an estimation point to capture relevant measurements. This strategy is effective when there is a sufficient number of measurements around the estimation point [15]; however, it performs poorly in sparsely sampled regions (e.g., at the start of a mission or near the boundary of the field).

To address these challenges, an adaptive method of measurement truncation is proposed wherein a standard rectangular geometric selector is used if it contains a threshold $M_{\min}$ number of measurements, and otherwise a nearest-neighbor selector is used to guarantee a minimum number of measurements (i.e., by

considering measurements outside the geometric selector, if needed). The geometric selector is denoted $G_R(\mathbf{x}, \mathbf{X}; w_G, h_G) = \{\mathbf{x}_i \in \mathbf{X} \mid \mathbf{x}_i \in R(\mathbf{x}; w_G, h_G)\}$ for all $i = \{1, 2, \dots, M\}$, where $R(\mathbf{x}; w, h) \subset \mathbb{R}^2$ denotes a rectangular area centered on $\mathbf{x} = (s_i, s_j) \in \mathbb{R}^2$ with width w and height h . The nearest-neighbor selector is denoted $G_{NN}(\mathbf{x}, \mathbf{X}; M_{\min}) \subset \mathbf{X}$ and it selects the subset of at most $M_{\min}$ measurement locations that are nearest-neighbors to $\mathbf{x}$ [53]. The adaptive selector is then

$$G(\mathbf{x}, \mathbf{X}) = \begin{cases} G_R(\mathbf{x}, \mathbf{X}; w_G, h_G) & \text{if } |G_R(\mathbf{x})| > M_{\min} \\ G_{NN}(\mathbf{x}, \mathbf{X}; M_{\min}) & \text{if } |G_R(\mathbf{x})| \leq M_{\min} \end{cases}, \quad (29)$$

where $|G_R(\mathbf{x})|$ is the number of measurements in the geometric selector.

4.3 Common-Data-Neighborhoods (CDNs)

Spatial truncation reduces the size of matrices required for inversion during the estimation process. However, it still requires inverting unique semivariogram matrices for every point to be estimated (since the geometric selector moves with the estimation point). Given many points of interest, the computation time of inverting many truncated matrices may exceed inverting the original (full measurement set) semivariogram. To reduce the number of required matrix inversions for kriging we modify the common data neighborhood (CDN) approach proposed in [48]. A group of estimation points are assigned a common semivariogram matrix based on nearby measurements (i.e., a fixed adaptive selector is used for multiple nearby estimation points). This approach is implemented as follows. The set of common data neighborhoods, $\mathcal{D} = \{\mathcal{D}_1, \dots, \mathcal{D}_{N_D}\}$, is defined as a collection of N_D disjoint rectangular regions whose union covers the entire survey area Q . Each data neighborhood $\mathcal{D}_i = R(\mathbf{n}_i; w_{\mathcal{D}}, h_{\mathcal{D}})$ is centered at a point $\mathbf{n}_i \in \mathbb{R}^2$ and is defined by a width $w_{\mathcal{D}}$ and height $h_{\mathcal{D}}$. Let $G_i = G(\mathbf{n}_i, \mathbf{X}; w_G, h_G, M_{\min})$ be the adaptive selector (29) for data neighborhood $\mathcal{D}_i$, with $w_G \geq w_{\mathcal{D}}$ and $h_G \geq h_{\mathcal{D}}$.

4.4 Heterogeneous CDN Kriging

The proposed estimation approach combines HFK specialized to heterogeneous robots, adaptive selectors, and CDNs. Suppose that an estimation point lies with the i th CDN, $\mathbf{x}_0 \in \mathcal{D}_i$. For all such estimation points, the HFK estimator (12) is used with the modifications (26)–(27) and a subset of measurements $\mathbf{X}$ given by the adaptive selector (29). Importantly, this same truncated data set is used for all estimation points in $\mathcal{D}_i$




Figure 3: Example of the proposed adaptive truncation strategy with $M_{\min} = 5$. Black circles are grid points, and blue ‘x’s are measurement locations. Two CDNs $\mathcal{D}_i$ and $\mathcal{D}_j$, centered on grid points $\mathbf{p}_i$ and $\mathbf{p}_j$, respectively, are highlighted in purple. When estimating a point A in $\mathcal{D}_i$ the corresponding rectangular selector contains more than $M_{\min}$ measurements and these measurements are used to define $\tilde{\mathbf{\Gamma}}$ in (12). When estimating a point B in $\mathcal{D}_j$ the nearest-neighbor selector is used to define $\tilde{\mathbf{\Gamma}}$ instead since insufficient measurements are located within the corresponding rectangular selector.

and hence the inversion in (12) need only occur once per neighborhood. The kriging weights and Lagrange multiplier for the point $\mathbf{x}_0 \in \mathcal{D}_i$ are

$$\begin{bmatrix} \lambda_{\text{HC}} \\ \mu_{\text{HC}} \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{\Gamma}}(G(\mathbf{n}_i, \mathbf{X})) & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix}^{-1} \begin{bmatrix} \tilde{\gamma}(G(\mathbf{n}_i, \mathbf{X}), \mathbf{x}_0) \\ 1 \end{bmatrix}, \quad (30)$$

and the estimate and variance are denoted with a subscript HC,

$$\hat{Z}_{\text{HC}}(\mathbf{x}_0) = \lambda_{\text{HC}}^T \tilde{\mathbf{Z}}(G(\mathbf{n}_i, \mathbf{X})) \quad (31)$$

and

$$\hat{\sigma}_{\text{HC}}(\mathbf{x}_0) = \tilde{\gamma}^T \lambda_{\text{HC}}(G(\mathbf{n}_i, \mathbf{X}), \mathbf{x}_0) + \mu_{\text{HC}}. \quad (32)$$

This approach is amenable to parallelization and permits a trade-off between computational efficiency and accuracy by adjusting the minimum number of considered measurements, $M_{\min}$, the dimensions of the common data/search neighborhoods, $w_{\mathcal{D}}, h_{\mathcal{D}}, w_G, h_G$.

4.5 Numerical Comparison

The heterogeneous, common-data-neighborhood kriging (HC) estimator was compared through numerical simulations to a HFK estimation [49] (without adaptive selectors or CDNs) and to ordinary kriging. The simulations consisted of generating a GP within a normalized domain $Q = [0, 1]^2$ using hyper-parameters $\boldsymbol{\theta} = [\zeta, \omega, \sigma_0^2]^T = [0, 0.3, 1]^T$. Measurement locations were randomly selected and the measurements were polluted with noise according to (9). For half

of the measurements, $\sigma_{\eta,1}^2 = 0.1$, for the other half, $\sigma_{\eta,2}^2 = 0.5$ to represent data collection by heterogeneous robots with different quality sensors. The noisy measurements were processed by the HC estimator with $h_{\mathcal{D}} = w_{\mathcal{D}} = 0.5\omega$, $h_G = w_G = \omega$ and $M_{\min} = 20$. The same measurements were also used to estimate the field with HFK and OK using a naive approach for nugget selection: $\zeta = (0.1 + 0.5)/2$.

The estimate was computed over one hundred unique realizations of the GP. For each realization, each method (HFK, HC, ordinary kriging) was used with 500, 1000, and 2000 measurements and the mean ME (24) and computation time were recorded. The results are shown in Fig. 4. An approximately ten-fold computation time savings was achieved when comparing HC to HFK and ordinary kriging. Also, HFK with heterogeneous measurement noise variance is superior to the naive approach of adjusting the nugget when using ordinary kriging. The accuracy reduction with HC compared to HFK is minimal. Note that accuracy did not improve with increased number of measurements for ordinary kriging—this is due to numerical instabilities in inverting large matrices, an effect which is pronounced where measurement noise is not considered.




Figure 4: A comparison between mapping error and computation time for three kriging methods. Error bars indicate one standard deviation.

5 Adaptive Sampling Algorithm

Since the actual realization of the field $Z(\mathbf{x})$ is not known to the robots, the mapping error (24) cannot be computed online for path planning. Moreover, online optimization of (25) for adaptive sampling is intractable due to high-dimensionality of the waypoint decision space and the nonlinear and differential constraints. Instead, we propose to use a surface of sam-

pling priority

$$\mathcal{J}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p}, \alpha) = [\alpha + \sigma^2(\mathbf{X}; \mathbf{p})] \cdot |\hat{\mathbf{Z}}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p})|, \quad (33)$$

to guide the allocation of waypoints. Equation (33) is evaluated over the grid of points $\mathbf{p} \in \mathcal{P}$ to give a time-varying surface that depends on the measurements $\tilde{\mathbf{Z}}$ and corresponding sample locations $\mathbf{X}$ that accumulate during the mission. We hypothesize that by allocating waypoints where sampling priority (33) is high a low cost (24) can be achieved. The sampling priority (33) includes a constant parameter $\alpha \geq 0$ and is a function of the uncertainty σ^2 and the estimated value $\hat{\mathbf{Z}}$ at a point $\mathbf{p}$ computed using a GP regression, such as (31)–(32). The term $[\alpha + \sigma^2(\mathbf{X}; \mathbf{p})]$ assigns a high priority to unexplored areas where the uncertainty is large. Multiplying by the estimate $|\hat{\mathbf{Z}}|$ scales the priority in areas that are uncertain and are predicted to deviate significantly from the field mean (i.e., peaks or valleys). Equation (33) assumes a zero-mean process but it can be modified to account for a known or estimated mean. The constant parameter α prevents the point-wise priority from reaching zero in regions that have already been sampled. Thus, α can be used as a tuning parameter to balance exploring new regions with returning to previously visited locations.

5.1 Algorithm Overview

The adaptive sampling algorithm proposed extends the Voronoi-based path generation approach in [12] to consider robots with heterogeneous sensing, heterogeneous dynamics, and mission time constraints. As the mission proceeds, a centroidal Voronoi tessellation algorithm periodically partitions the time-varying sampling priority surface to identify high-value sampling locations. The path for each robot is modeled as an mechanical system: a sequence of masses (waypoints) are interconnected by springs and dampers and pulled towards Voronoi cell centroids (see Fig. 5). At each path planning cycle the robots are iteratively simulated following their respective waypoint paths as stiffness/damping parameters are adjusted to satisfy mission time constraints under the robots' heterogeneous dynamics. A detailed implementation of the algorithm is described below.

5.2 Algorithm Detailed Implementation

Algorithm 1 begins with the initialization of $N \cdot N_w$ unique waypoints $\mathbf{W}$ within the sample space $\mathcal{Q}$ [A1.2] (e.g., using a nominal lawnmower coverage path for each robot). This waypoint set is then discretized into




Figure 5: A constrained CVT of three waypoints visualized as gray circles with mass m and associated capture radius. Voronoi cells are defined by (17). Centroids are depicted as blue circles. The time-varying sampling priority surface is indicated by the contour plot with warmer colors representing regions of higher priority. Pairs of adjacent waypoints are connected by a spring with stiffness k_p and a damper with damping coefficient b . Waypoints are connected to their associated centroids with spring constant, k_c , and a damper. A dynamically feasible trajectory through the waypoints set is shown in blue.

N_c planning cycles consisting of $C = \lceil N_w/N_c \rceil$ waypoints per cycle [A1.3], where $\lceil \cdot \rceil$ denotes the ceiling operator. The integer index g denotes the current cycle. At each new planning cycle, recently acquired measurements $\tilde{\mathbf{Z}}$ at locations $\mathbf{X}$ are assimilated via a GP regression (Sec. 4) from which the sampling priority surface (33) is evaluated point-wise (at each grid point in $\mathcal{P}$) [A1.5] and arranged as a matrix

$$\mathfrak{J} = \begin{bmatrix} \mathcal{J}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p}_{1,1}, \alpha) & \dots & \mathcal{J}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p}_{1,n}, \alpha) \\ \vdots & \ddots & \vdots \\ \mathcal{J}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p}_{m,1}, \alpha) & \dots & \mathcal{J}(\mathbf{X}, \tilde{\mathbf{Z}}; \mathbf{p}_{m,n}, \alpha) \end{bmatrix}, \quad (34)$$

where $\mathbf{p}_{i,j}$ is the grid point associated with the i th row and the j th column. For the first iteration, when no measurements have been performed, the priority at all grid points is initialized to a non-zero constant. The sampling priority matrix $\mathfrak{J}$ is then passed, along with the initialized waypoints $\mathbf{W}$, to a modified centroidal Voronoi path generation algorithm, **mCVPG** [A.2].

The **mCVPG** [A.2] algorithm select high-priority waypoints without considering the differential constraints and mission time. The supplied initial waypoints $\mathbf{W}$ are used as generating points, and **mCVPG** recursively calculates the Voronoi tessellation [A2.3] to determine Voronoi cell centroids (using $\mathfrak{J}$) until a convergence criteria is met. At each iteration the waypoints $\mathbf{W}$ are moved according to the dynamics of a mass-spring-damper network. Waypoints are treated as having equal mass m and are connected by springs to their corresponding Voronoi cell centroid and to adjacent waypoints in the path for each agent. Dampers are

Algorithm 1 Satisfying mission time constraint using mCVPG with multiple robots

Require: 1
Require: $T_\Delta \geq 0$ // mission time tolerance
Require: N // number of robots
Require: N_w // total num. waypoints in mission
Require: N_c // number of planning cycles
Require: γ_0 // init. spring, damp. constants
Require: $Q \subset \mathbb{R}^2$ // survey area
Require: $\mathcal{P}$ // evaluation grid points set
Require: I // maximum iterations
Require: I_s // maximum iterations for mCVPG
Require: $\{\bar{v}_i\}_{i=1}^N$ // robots' target speed
1: $\mathbf{X} \leftarrow \emptyset, \tilde{\mathbf{Z}} \leftarrow \emptyset$ // initialize samples
2: $\mathbf{W} \leftarrow \text{initializeWaypoints}(N, N_w, Q)$
3: $C \leftarrow \lceil N_w / N_c \rceil$ // num. wpts per cycle
4: **for** $g = \{1, 2, \dots, N_c\}$ **do**
5: $\mathfrak{J} \leftarrow \text{evaluateCost}(\mathbf{X}, \tilde{\mathbf{Z}}, \mathcal{P})$
6: $\{\gamma_i\}_{i=1}^N \leftarrow \text{initSDs}(\gamma_0)$
7: $b \leftarrow 0$ // iteration counter
8: **while** $b < I$ **do**
9: $\mathbf{W} \leftarrow \text{mCVPG}(\mathbf{W}; N, \{\gamma_i\}_{i=1}^N, g, C, \mathfrak{J}, 10^{-4}, I_s)$
10: **for** $i = 1, \dots, N$ **do**
11: $t_i \leftarrow \text{simulateAgent}(\mathbf{W}_i; \mathbf{s}_i(T(\mathbf{w}_{i,gC})), g, C, \bar{v}_i)$
12: $\epsilon_i \leftarrow T_m - t_i$ // sim. time error
13: **if** $|\epsilon_i| > T_\Delta$ **then**
14: $\gamma_i \leftarrow \text{updateSDs}(\epsilon_i, T_m, \gamma_i)$
15: **end if**
16: **end for**
17: **if** $|\epsilon_i| \leq T_\Delta$ for all $i = 1, \dots, N$ **then**
18: **break**
19: **end if**
20: $b \leftarrow b + 1$
21: **end while**
22: $(\mathbf{X}, \tilde{\mathbf{Z}}, t_m) \leftarrow \text{collectData}(\mathbf{W}, N, g, C, \mathbf{X}, \tilde{\mathbf{Z}})$
23: $T_m \leftarrow T_m - t_m$ // remaining mission time
24: **end for**

placed alongside each spring to model energy dissipation as the system comes to a rest at an equilibrium position. The values of the spring constants, dampers, and relaxed spring length are uniform across a robot's path but differ between robots, $\gamma_i = \{k_{p,i}, k_{c,i}, b_i, d_i\}$ for $i = 1, \dots, N$. The set of configuration parameters is initialized at the beginning of each planning cycle with user-defined values in `initSDs` [A1.6]. Figure 5 visualizes this spring-mass-damper system for a series of three waypoints. Each waypoint, $\mathbf{w}_{i,j}$, excluding the first and last elements, has three linear spring and damper connections: one to its previous neighbor $\mathbf{w}_{i,j-1}$, one to its next neighbor ($\mathbf{w}_{i,j+1}$), and one to its corresponding centroid ($\mathbf{c}_{V_{i,j}}$). The spring force exerted upon the j th waypoint by its adjacent waypoints is:

$$\mathbf{f}_{i,j}^{(p)} = k_{p,i}[(\|\mathbf{r}_{i,j-1}\| - d_i)\hat{\mathbf{r}}_{i,j-1} + (\|\mathbf{r}_{i,j+1}\| - d_i)\hat{\mathbf{r}}_{i,j+1}], \quad (35)$$

where $\mathbf{r}_{i,j\pm 1} = \mathbf{w}_{i,j} - \mathbf{w}_{i,j\pm 1}$ is the vector distance between waypoint j and the next waypoint, $(j+1)$, or the previous waypoint, $(j-1)$. The normalized dis-

Algorithm 2 Multi-robot centroidal Voronoi path generation (mCVPG)

Require: $\{\mathbf{W}_i\}_{i=1}^N$ // waypoints
Require: N // number of robots
Require: $\{\gamma_i\}_{i=1}^N$ // spring, damp. constants
Require: g // current cycle
Require: C // number of waypoints per cycle
Require: $\mathfrak{J}_k$ // sampling priority surface
Require: η // stopping criteria
Require: I_s // maximum iterations
1: $b \leftarrow 0$ // iteration counter
2: **while** $b < I_s$ **do**
3: $V \leftarrow \text{voronoiTessellation}(\{\mathbf{W}_i\}_{i=1}^N)$
4: $\mathbf{c}_V \leftarrow \text{voronoiMassCentroid}(V, \mathfrak{J}_k)$
5: **for** $i = 1, \dots, N$ **do**
6: $\mathbf{f}_i \leftarrow \text{calcForce}(\mathbf{W}_i, \gamma_i, \mathbf{c}_V)$
7: $(\mathbf{W}_i, \tilde{\mathbf{W}}_i) \leftarrow \text{movePts}(\mathbf{W}_i, \mathbf{f}_i, g, C)$
8: **end for**
9: **if** $\|\dot{\mathbf{w}}_{i,j}\| \leq \eta$ for all $i = 1, \dots, N$ and $j = 1, \dots, N_w$ **then**
10: **break** // waypoints converged
11: **end if**
12: $b \leftarrow b + 1$
13: **end while**
14: **return** $\{\mathbf{W}_i\}_{i=1}^N$ // updated waypoints

tance vector is $\hat{\mathbf{r}}_{i,p} = \mathbf{r}_{i,p} / \|\mathbf{r}_{i,p}\|$, d_i is the spring's unstretched length, and $k_{p,i}$ is a spring constant. The relaxed spring length is updated prior to each planning cycle to be $d_i = 0.5\bar{v}_i(T_m - T(\mathbf{w}_{i,gC})) / (N_w - gC)$, where $\bar{v}_i$ is the i th robot's target speed. This relaxed spring length heuristic divides the expected distance robot i will travel in the remaining time $T_m - T(\mathbf{w}_{i,gC})$ by the number of remaining waypoints. The remaining $N_w - gC$ are then equally spaced along the remaining path. The force exerted on the j th waypoint by its centroid is $\mathbf{f}_{i,j}^{(c)} = k_{c,i}(\mathbf{c}_{V_{i,j}} - \mathbf{w}_{i,j})$ where $k_{c,i}$ is a nominal constant for the centroidal attracting spring force stiffness. Additionally, the linear damping force is $\mathbf{f}_{i,j}^{(b)} = -b_i\dot{\mathbf{w}}_{i,j}$ where $\mathbf{w}_{i,j}$ is the velocity of the waypoint. The total force [A2.6] acting on $\mathbf{w}_{i,j}$ is

$$\mathbf{f}_{i,j} = \mathbf{f}_{i,j}^{(p)} + \mathbf{f}_{i,j}^{(c)} + \mathbf{f}_{i,j}^{(b)}. \quad (36)$$

The equation of motion for waypoint $\mathbf{w}_{i,j}$ with mass m is then $\ddot{\mathbf{w}}_{i,j} = \mathbf{f}_{i,j}/m$. The function `movePts` performs Euler integration to determine the new position of each waypoint given the associated total force (36). Waypoints that have already been visited during the mission no longer need updating. Thus, if a waypoint is associated with a future cycle, i.e., the waypoint index is greater than gC , the position is updated and the velocity is recorded. Otherwise, the waypoint is frozen and the associated velocity is zero [A2.7]. The system is numerically integrated until the speed of each waypoint is below a desired threshold, η [A2.9] or a maximum number of iterations is reached [A2.2].

Once mCVPG has returned an updated set of waypoints each robot is simulated moving through its remaining waypoints from its current state [A1.11] and the estimated time to traverse all remaining waypoints is computed. This simulation time, t_i , is compared with the remaining mission time [A1.13]. If the difference exceeds the user defined mission time tolerance, T_Δ , the spring-mass-damper system parameters are updated (i.e., stiffness between waypoints and Voronoi centroids is increased when the simulated mission time is less than required and decreased otherwise). In practice, a buffer $\mathbf{E}$ records the mission time error from previous iterations for each robot [A1.12]. Let $\max(\mathbf{E}_i)$ be the maximum mission time error for the i th robot. Similar to [12], the spring stiffness is updated for the next optimization iteration as

$$k_{c,i} \leftarrow \begin{cases} k_{c,i} \left(\frac{\epsilon_i}{\max(\mathbf{E}_i)} + 1 \right) & \text{if } t_i + T_\Delta < T_m \\ k_{c,i} \left(\frac{\epsilon_i}{\max(\mathbf{E}_i)} + 1 \right)^{-1} & \text{if } t_i - T_\Delta > T_m \\ k_{c,i} & \text{if } |t_i - T_m| \leq T_\Delta \end{cases} \quad (37)$$

and the damping coefficient for the i th robot as $b_i = 0.5\sqrt{\max(k_{p,i}, k_{c,i})}$ [A1.14]. The optimization loop [A1.8–21] runs, calculating a new constrained CVT solution with updated parameters during each optimization step, until the simulated mission time for all robots is met. A maximum number of iterations of this outer loop optimization is enforced [A1.8]. Once the optimized waypoints are found, the robots execute them and the process repeats at the next planning cycle. To account for any discrepancies between simulated and actual mission time (relevant during real-world experiments), the time elapsed since the previous planning cycle is reported and the remaining mission time is updated [A1.23].

6 Simulation Study

A Monte Carlo simulation was designed to compare performance of the adaptive sampling algorithm to a non-adaptive lawnmower survey. The Monte Carlo simulation tested 30 unique GP realization for each combination of parameters listed in Table 2, resulting in 3,600 simulations. Each GP realization was constructed using $\boldsymbol{\theta} = [\zeta, \omega, \sigma_0^2]^T = [0, \omega_t, 1]^T$ as the hyper-parameters where ω_t is a selected entry from Table 2. The HC estimation parameters used were: $w_D = h_D = 0.5\omega$, $w_G = h_G = 1.5\omega$ and $M_{\min} = 10$.

Each trial compared mapping error (24) for an adaptive and a non-adaptive/control pair of robots. First, the path of the control robots was planned by designating one of them the leader and assigned them

Parameter	Symbol	Value
Number of robots	N	2
Agent speed	v	$\{0.01, 0.02, 0.04\}$ LU/s
Measurement noise	σ_η^2	$\{0.0, 0.05, 0.15\}$
GP length-scale	ω_t	$\{0.1, 0.2, 0.3, 0.5\}$ LU
No. swaths for leader robot	N_l	$\{6, 7, 8, 9, 10\}$
No. planning horizons	N_c	10

Table 2: Parameters used in Monte Carlo simulation. Simulations were conducted using a length unit (LU) that resulted in a survey area with unit width and height.




Figure 6: Example simulation scenario featuring a leader and follower robot. The leader is assigned four swaths and the follower is assigned three. Paths are positioned within the field for uniform coverage. The dark shaded regions depict already surveyed regions of the field.

a target speed and number of swaths (see Fig. 6). The waypoints were defined in terms of swath width, $c_w = Q_w / (N_l + 1)$ where Q_w is the width of the survey field and N_l is the total number of swaths. The points were then inset from the border by a distance $c_w/2$. The leader was simulated following its assigned waypoints and the total mission time was recorded. The follower robot's target speed was then iteratively adjusted to ensure that it achieved its final waypoint at the same time as the leader. Once mapping error for the control robots was quantified the adaptive robots were simulated with an identical GP realization for that trial and mapping error was recorded.

Figure 7 shows changes in mapping error for robots with equal speed (0.1 LU/s) and for GPs with increasing length-scale for both noisy and noise-free measurement cases. In the case of noise-free measurements, the adaptive algorithm exceeds the performance of the lawnmower pattern in fields where the GP's length-scale is less than 50% of the the survey area's length. These corresponds to GP realizations with a high degree of variability. As the GP length-scale approaches 50% of the survey area's length, the adaptive algorithm and lawnmower patterns converge to similar mapping errors given that there is little variation in the process to observe in the survey area. When comparing the noisy measurements, the same trend exists but is

less pronounced. The smoothing effect of measurement error filtered kriging adds difficulty to the estimation process which makes path selection less efficient.




Figure 7: Monte Carlo simulation results given parameters specified in Table 2 comparing final mean mapping error over all simulated length-scales when both robots’ speeds are 0.01 LU/s and given no measurement noise variance (noise-free measurements) and $\sigma_{1,\eta}^2 = 0.15, \sigma_{2,\eta}^2 = 0.05$ for measurement error case (noisy measurements). Error bars indicate one standard deviation.

Table 3 compares the simulation results for a small GP length-scale ($\omega = 0.1$) under different combinations of heterogeneous dynamics and sensing. In all cases the adaptive algorithm outperforms the non-adaptive lawnmower strategy, although the advantage is greatest in the case of heterogeneous dynamics with equal sensing capabilities. Considering the error bars of the simulation, the results suggest that the proposed approach leads to statistically significant improvements in performance by effectively leveraging heterogeneous sensing and heterogeneous dynamics.

Lastly, Fig. 8 illustrates the multi-robot sampling trajectories that result from adaptive sampling algorithm with heterogeneous robots. Figures 8a–8d are the completed and planned trajectories and sampling priority surface at different stages during the mission.

7 Experimental Demonstration

The adaptive sampling algorithm was deployed on a pair of autonomous surface vessels (ASVs) mapping bathymetry in a small section of a lake. This section describes the ASV platforms, the collection of ground truth data, and the experimental results.

7.1 Autonomous Surface Vessels

Two custom-built autonomous surface vessels were used in the experiments. The control framework for each ASV was run on an Intel NUC (10i7FNH) with

a 4.7GHz i7 CPU and 32GB of RAM. Communication between each robot and a ground station (Dell Vostro Notebook 7500) used Wi-Fi and a 900 MHz radio (RFD900+ modem) for communication. Each robot was outfitted with a BlueRobotics Ping Sonar capable of measuring depths up to 70 m. A pair of 14.8 V and 6000 mAh LiPo batteries provided up to one hour of mission time per ASV. The mechanical, electrical, and software design is detailed in [54].

7.2 Ground Truth Survey

Prior to adaptive sampling experiments, a survey was conducted in an approximately 140 m x 140 m small rectangular section of Lake Norman, NC to generate a ground truth map of the bathymetry. The survey consisted of lawnmower pattern with waypoints arranged to form eight swaths. Approximately 5,000 data points were collected during the survey and a semivariogram was fit to the binned data, as shown in Fig. 9. The curve fit was used to determine the hyperparameters $\theta = [\zeta, \omega, \sigma_0^2]^T = [0, 153 \text{ m}, 7.27 \text{ m}^2]^T$.

7.3 Experimental Demonstration

The ability of our algorithm to improve performance under heterogeneous sensing and mobility was illustrated in Sec.6 through simulation results. For the experiment, the team was limited to using available equipment consisting of identical ASVs with the same onboard sensors and propulsion system (see Fig. 10a). The purpose of the experiments was therefore to demonstrate the technical feasibility of implementing the algorithms onboard a real system and report lessons learned in doing so, rather than to quantifying performance under heterogeneity. Moreover, because of the limited variability of the bathymetry in the survey area (i.e., a large length scale similar to the rightmost comparison in Fig. 7) we expect the adaptive and non-adaptive algorithms to produce similar mapping error.

Experimental trials were conducted in the previously surveyed operating region described in Sec. 7.2. The adaptive sampling algorithm used the parameters listed in Table 4. Several trials were conducted, each lasting between 7-10 minutes. This paper discusses one of the trials; further experimental details can be found in [54].

Through testing the team discovered that several modifications were needed to the estimation and adaptive sampling algorithms. To account for the computation time required during planning, the re-planning cycle was modified to start early within each cycle. That

Cases	Dynamics		Sensing		Results
	Equal	Hetero.	Equal	Hetero.	
1	X		X		1.16
2		X	X		1.41
3	X			X	1.05
4		X		X	1.08

Table 3: Comparison of simulation results given all combinations of dynamics and measurement error variance. The bar graph indicates ratio of the mapping error of the adaptive sampling mission to the control sampling mission. The red line represents parity, a ratio of 1:1. Simulations performed over a scalar field with $\sigma_0^2 = 1$ and $\omega = 0.1$. Ninety simulations split evenly among the discrete mission classifications.




Figure 8: Examples of adaptive patterns given $v=0.01$ LU/s, $N_l = 3$, and $\omega = 0.2$ given measurement noise $\sigma_{1,\eta}^2 = 0.15$ for the red robot and $\sigma_{2,\eta}^2 = 0.05$ for the white robot. Panels (a)-(d) show the sampling priority surface and planned paths at different stages of the mission along with the sampling priority surface. Panel (c) is the path of the control robots and Panel (f) is the actual field in this example.

is, rather than waiting to complete all waypoints in the current cycle re-planning started p_r waypoints prior to the next cycle. The first robot that reached this target triggered the computation on the ground station. Although this strategy ignores some measurements acquired near the end of each planning cycle, it allows the ASVs to smoothly transition to updated trajectories. To compensate for a non-zero p_r , the waypoints remaining in the current cycle are left unmodified along with all previously reached waypoints in the `movePts` function [A1.7].

Given bandwidth limitations with multi-robot communication, the sensing interval, T_s was set at 10 seconds. The relatively low number of measurements justified the use of HFK estimation rather than HC estimation. The measurement noise used for both robots with HFK was $\sigma_\eta^2(\mathbf{x}) = 0.005d(\mathbf{x})$ where $d(\mathbf{x})$ is the depth measured at spatial location $\mathbf{x}$ based upon the manufacturer's specifications [55].

The trajectories of the ASVs during the experiment are shown in Fig. 10b. The ASVs were launched from a nearby dock and traveled to their first waypoint (in-




Figure 9: Experimental semivariogram computed using data gathered during ground truth survey on Lake Norman, NC.

Parameter	Symbol	Value
Number of cycles	N_c	4
Number of waypoints	N_w	32
Number of waypoints per cycle	C	8
Early re-planning offset	p_r	2
Sensing interval	T_s	10 s
Measurement noise variance	σ_η^2	Variable ^[1]
Mission time tolerance	T_Δ	± 10 s

Table 4: Common mission parameters among missions. [1] Site-specific measurement noise variance is 0.5% of measured depth [55]

icated by a red marker). The mission time allowed for three additional planning cycles. Figure 10c depicts the samples that were collected and the resulting kriging estimate computed using the estimated semivariogram parameters (Fig. 9). The result compares favorably to the actual bathymetry determined during the ground truth survey (Fig. 10d) and to publicly available bathymetric maps of the corresponding region. During the experiment, the path of the left-most agent deforms from the nominal lawnmower pattern (used to initialize the algorithm) at two locations: near the valley located at $(x, y) = (35 \text{ m}, 100 \text{ m})$ and towards the peak located at $(x, y) = (120 \text{ m}, 140 \text{ m})$. The first upward swath of the right-most agent maintains a near constant spacing to the downward swath of the left-most agent and is also pulled toward the peak located at $(x, y) = (120 \text{ m}, 140 \text{ m})$. During the second half of the mission the right-most agent allocates sampling effort in the central region of the space where the gradients are large. The fact that the adaptive algorithm looks similar to a lawnmower with small deviations at a few points is not surprising and consistent with the simulation results that show the adaptive algorithm and non-adaptive lawnmower perform similarly when survey area has limited variability. The adaptive sampling algorithm is expected to be most effective in regions

that exhibit more variation. A survey over a much larger portion of the lake with longer endurance/range vehicles outfitted with different sensing and propulsion equipment would be more favorable for illustrating our approach experimentally.

8 Conclusion

This paper presented an approach for adaptive sampling of a stationary, isotropic, Gaussian process with a mobile sensor network consisting of robots with heterogeneous dynamics and varying measurement noise variance. A Gaussian process regression was proposed to handle the heterogeneous noise variance and use common data neighborhood to reduce computational complexity. This estimation procedure is combined with an adaptive sampling algorithm to periodically replan the trajectory of each robot to visit high priority sampling regions and meet mission time constraints. The proposed strategy was evaluated through a numerical Monte Carlo simulations and real-world experiments. The numerical simulations showed that the adaptive sampling strategy substantially reduced mapping error in comparison to a non-adaptive lawnmower survey. The feasibility of the approach was demonstrated by deployed the algorithm on a team of two ASVs to survey a small section of a lake.


Future work should consider optimizing the adaptive sampling algorithm to allow for efficient scaling for a large number of robots in a decentralized framework, sampling time-varying spatial processes, estimation of hyperparameters on-the-fly, and/or simultaneous estimation of multiple spatial attributes that are statistically correlated (i.e., cokriging with heterogeneous measurement types). Outdoor testing with different ASVs can allow for more extensive experimental evaluation.

References


- [1] D. A. Paley and A. Wolek, “Mobile sensor networks and control: Adaptive sampling of spatiotemporal processes,” *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 91–114, 2020.
- [2] S. Bai, T. Shan, F. Chen, L. Liu, and B. Englot, “Information-driven path planning,” *Current Robotics Reports*, vol. 2, pp. 177–188, 2021.
- [3] J. Pulido Fentanes, A. Badiee, T. Duckett, J. Evans, S. Pearson, and G. Cielniak, “Kriging-based robotic exploration for soil moisture map-




(a) Autonomous surface vessels



(b) Path planning cycles



(c) Estimated bathymetry



(d) Actual bathymetry

Figure 10: Field experiment conducted in Lake Norman, NC. (a) The experiment used the two ASVs with a mission time $T_m = 430$ s and a speed of $v = 0.7$ m/s to map the operating region indicated by the white box. (b) The beginning of each replanning cycle is shown as a colored waypoint for a total of four cycles. (c) The samples used to estimate the bathymetry are shown as black markers plotted with the predicted mean over the operating region. (d) The actual bathymetry determined from a ground truth survey.

ping using a cosmic-ray sensor,” *Journal of Field Robotics*, vol. 37, no. 1, pp. 122–136, 2020.

- [4] K.-C. Ma, L. Liu, H. K. Heidarrsson, and G. S. Sukhatme, “Data-driven learning and planning for environmental sampling,” *Journal of Field Robotics*, vol. 35, no. 5, pp. 643–661, 2018.
- [5] T. O. Fossum, J. Eldsvik, I. Ellingsen, M. O. Alver, G. M. Fragoso, G. Johnsen, R. Mendes, M. Ludvigsen, and K. Rajan, “Information-driven robotic sampling in the coastal ocean,” *Journal of Field Robotics*, vol. 35, pp. 1101–1121, 2018.
- [6] N. E. Leonard, D. A. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. E. Davis, “Collective motion, sensor networks, and ocean sampling,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 48–74, 2007.
- [7] R. Marchant and F. Ramos, “Bayesian optimisation for intelligent environmental monitoring,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 2242–2249.
- [8] R. A. Olea, *Geostatistics for Engineers and Earth Scientists*. New York, USA: Springer Science+Business Media, 2003.
- [9] 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.
- [10] 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 International Conference on Robotics and Automation*, 2018, pp. 6220–6227.
- [11] M. Jadaliha, Y. Xu, J. Choi, N. S. Johnson, and W. Li, “Gaussian process regression for sensor networks under localization uncertainty,” *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 223–237, 2012.
- [12] K. R. Jensen-Nau, T. Hermans, and K. K. Leang, “Near-optimal area-coverage path planning of energy-constrained aerial robots with application in autonomous environmental monitoring,” *IEEE Transactions on Automation Science and Engineering*, vol. 18, pp. 1453–1468, 2021.
- [13] A. Singh, F. Ramos, H. D. Whyte, and W. J. Kaiser, “Modeling and decision making in spatio-temporal processes for environmental surveillance,” in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 5490–5497.
- [14] G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [15] Y. Xu, J. Choi, and S. Oh, “Mobile sensor network navigation using Gaussian processes with truncated observations,” *IEEE Transactions on Robotics*, vol. 27, no. 6, pp. 1118–1131, 2011.
- [16] M. Jadaliha and J. Choi, “Environmental monitoring using autonomous aquatic robots: Sampling algorithms and experiments,” *IEEE Transactions on Control Systems Technology*, vol. 21, no. 3, pp. 899–905, 2013.
- [17] L. V. Nguyen, S. Kodagoda, R. Ranasinghe, and G. Dissanayake, “Information-driven adaptive sampling strategy for mobile robotic wireless sensor network,” *IEEE Transactions on Control Systems Technology*, vol. 24, no. 1, pp. 372–379, 2015.

- [18] D. Gu and H. Hu, "Spatial gaussian process regression with mobile sensor networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 23, no. 8, pp. 1279–1290, 2012.
- [19] R. Graham and J. Cortés, "Adaptive information collection by robotic sensor networks for spatial estimation," *IEEE Transactions on Automatic Control*, vol. 57, no. 6, pp. 1404–1419, 2012.
- [20] C. Wang, L. Wei, Z. Wang, M. Song, and N. Mahmoudian, "Reinforcement learning-based multi-*auv* adaptive trajectory planning for under-ice field estimation," *Sensors*, vol. 18, no. 11, p. 3859, 2018.
- [21] S. Manjanna, M. A. Hsieh, and G. Dudek, "Scalable multirobot planning for informed spatial sampling," *Autonomous Robots*, vol. 46, no. 7, pp. 817–829, 2022.
- [22] K. Jakkala and S. Akella, "Multi-robot informative path planning from regression with sparse Gaussian processes," *arXiv preprint arXiv:2309.07050*, 2023.
- [23] G. Hitz, E. Galceran, M.-È. Garneau, F. Pomerleau, and R. Siegwart, "Adaptive continuous-space informative path planning for online environmental monitoring," *Journal of Field Robotics*, vol. 34, no. 8, pp. 1427–1449, 2017.
- [24] J. Zhang, M. Liu, S. Zhang, R. Zheng, and S. Dong, "Multi-*auv* adaptive path planning and cooperative sampling for ocean scalar field estimation," *IEEE Transactions on Instrumentation and Measurement*, vol. 71, pp. 1–14, 2022.
- [25] C. Xiao and J. Wachs, "Nonmyopic informative path planning based on global kriging variance minimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1768–1775, 2022.
- [26] W. Luo, C. Nam, G. Kantor, and K. Sycara, "Distributed environmental modeling and adaptive sampling for multi-robot sensor coverage," in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, 2019, pp. 1488–1496.
- [27] 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 International Conference on Automated Planning and Scheduling*, vol. 19, 2009, pp. 233–240.
- [28] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "Distributed robotic sensor networks: An information-theoretic approach," *The International Journal of Robotics Research*, vol. 31, no. 10, pp. 1134–1154, 2012.
- [29] G. P. Kontoudis and D. J. Stilwell, "Decentralized nested gaussian processes for multi-robot systems," in *2021 IEEE International Conference on Robotics and Automation*, 2021, pp. 8881–8887.
- [30] —, "Decentralized federated learning using gaussian processes," in *2023 International Symposium on Multi-Robot and Multi-Agent Systems*, 2023, pp. 1–7.
- [31] 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 International Conference on Intelligent Robots and Systems*, 2020, pp. 11 718–11 723.
- [32] M. Malencia, S. Manjanna, M. A. Hsieh, G. Pappas, and V. Kumar, "Adaptive sampling of latent phenomena using heterogeneous robot teams (ASLaP-HR)," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2022, pp. 8762–8769.
- [33] M. J. Ten Kathen, F. P. Samaniego, I. J. Flores, and D. G. Reina, "AquaHet-PSO: An informative path planner for a fleet of autonomous surface vehicles with heterogeneous sensing capabilities based on multi-objective PSO," *IEEE Access*, 2023.
- [34] J. Zhang, M. Liu, S. Zhang, R. Zheng, and S. Dong, "A path planning approach for multi-*auv* systems with concurrent stationary node access and adaptive sampling," *IEEE Robotics and Automation Letters*, vol. 9, no. 3, pp. 2343–2350, 2024.
- [35] S. McCammon, G. Marcon dos Santos, M. Frantz, T. P. Welch, G. Best, R. K. Shearman, J. D. Nash, J. A. Barth, J. A. Adams, and G. A. Hollinger, "Ocean front detection and tracking using a team of heterogeneous marine vehicles," *Journal of Field Robotics*, vol. 38, no. 6, pp. 854–881, 2021.
- [36] P. Tokekar, J. Vander Hook, D. Mulla, and V. Isler, "Sensor planning for a symbiotic UAV and UGV system for precision agriculture," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1498–1511, 2016.

- [37] 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 International Conference on Robotics and Automation*, 2021, pp. 8859–8865.
- [38] 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.
- [39] H. Qin, Z. Meng, W. Meng, X. Chen, H. Sun, F. Lin, and M. H. Ang, "Autonomous exploration and mapping system using heterogeneous uavs and ugvs in gps-denied environments," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 2, pp. 1339–1350, 2019.
- [40] T. Salam and M. A. Hsieh, "Heterogeneous robot teams for modeling and prediction of multiscale environmental processes," *Autonomous Robots*, vol. 47, no. 4, pp. 353–376, 2023.
- [41] R. Cui, Y. Li, and W. Yan, "Mutual information-based multi-AUV path planning for scalar field sampling using multidimensional RRT*," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 46, no. 7, pp. 993–1004, 2016.
- [42] J. Cortés, "Distributed Krige Kalman filter for spatial estimation," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2816–2827, 2009.
- [43] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. Boston, USA: MIT Press, 2006.
- [44] Y. Xu and J. Choi, "Spatial prediction with mobile sensor networks using Gaussian processes with built-in Gaussian Markov random fields," *Automatica*, vol. 48, no. 8, pp. 1735–1740, 2012.
- [45] H. Liu, Y.-S. Ong, X. Shen, and J. Cai, "When Gaussian process meets big data: A review of scalable GPs," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 11, pp. 4405–4423, 2020.
- [46] S. Agarwal and S. Akella, "Area coverage with multiple capacity-constrained robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3734–3741, 2022.
- [47] M. Wei and V. Isler, "Coverage path planning under the energy constraint," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 368–373.
- [48] M. Vignès, O. Kolbjørnsen, V. L. Hauge, P. Dahle, and P. Abrahamsen, "Fast and accurate approximation to kriging using common data neighborhoods," *Mathematical Geosciences*, vol. 49, pp. 619–634, 2017.
- [49] W. F. Christensen, "Filtered kriging for spatial data with heterogeneous measurement error variances," *Biometrics*, vol. 67, no. 3, pp. 947–957, 2011.
- [50] N. A. Cressie, *Statistics for Spatial Data*. New York, USA: Wiley-Interscience, 1993.
- [51] Q. Du, V. Faber, and M. Gunzburger, "Centroidal Voronoi tessellations: Applications and algorithms," *SIAM Review*, vol. 41, no. 4, pp. 637–676, 1999.
- [52] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129–137, 1982.
- [53] J. P. C. Kleijnen and W. C. M. van Beers, "Prediction for big data through kriging: Small sequential and one-shot designs," *American Journal of Mathematical and Management Sciences*, vol. 39, no. 3, pp. 199–213, 2020.
- [54] M. Brancato, "Autonomous sensing of a Gaussian spatial process with multiple heterogeneous agents," 2022, M.S. Thesis, Univ. North Carolina at Charlotte.
- [55] BlueRobotics, "Ping sonar altimeter and echosounder," <https://bluerobotics.com/store/sensors-sonars-cameras/sonar/ping-sonar-r2-rp/>, (accessed: 2022-12-30).