

Learning a Spatial Field in Minimum Time with a Team of Robots

Varun Suryan

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Computer Engineering

Pratap Tokekar, Chair
A. Lynn Abbott
Ryan K. Williams

August 12, 2019
Blacksburg, Virginia

Keywords: Sensor Placement, Informative Path Planning, Gaussian Processes
Copyright 2019, Varun Suryan

Learning a Spatial Field in Minimum Time with a Team of Robots

Varun Suryan

(ABSTRACT)

We study an informative path planning problem where the goal is to minimize the time required to learn a spatial field. Specifically, our goal is to ensure that the mean square error between the learned and actual fields is below a predefined value. We study three versions of the problem. In the placement version, the objective is to minimize the number of measurement locations. In the mobile robot version, we seek to minimize the total time required to visit and collect measurements from the measurement locations. A multi-robot version is studied as well where the objective is to minimize the time required by the last robot to return back to a common starting location called depot. By exploiting the properties of Gaussian Process regression, we present constant-factor approximation algorithms that ensure the required guarantees. In addition to the theoretical results, we also compare the empirical performance using a real-world dataset with other baseline strategies.

This material is based upon work supported by the National Science Foundation under Grant number 1637915 and NIFA grant 2018-67007-28380.

Learning a Spatial Field in Minimum Time with a Team of Robots

Varun Suryan

(GENERAL AUDIENCE ABSTRACT)

We solve the problem of measuring a physical phenomenon accurately using a team of robots in minimum time. Examples of such phenomena include the amount of nitrogen present in the soil within a farm and concentration of harmful chemicals in a water body etc. Knowing accurately the extent of such quantities is important for a variety of economic and environmental reasons. For example, knowing the content of various nutrients in the soil within a farm can help the farmers to improve the yield and reduce the application of fertilizers, the concentration of certain chemicals inside a water body may affect the marine life in various ways. In this thesis, we present several algorithms which can help robots to be deployed efficiently to quantify such phenomena accurately. Traditionally, robots had to be teleoperated. The algorithms proposed in this thesis enable robots to work more autonomously.

Dedication

To Family and Friends.

Acknowledgments

I am tremendously grateful to my advisor Pratap Tokekar and all the lab members. Research meetings with Dr. Tokekar have always been helpful and fun in preparing this thesis. Thanks to my committee members Lynn Abbott and Ryan Williams for the invaluable inputs.

Contents

List of Figures	vii
1 Introduction	1
2 Related Work and Background	5
2.1 Stationary Sensor Placement	5
2.2 Mobile Sensing	6
2.3 Sensing with Multiple Robots	8
2.4 Gaussian Processes	9
2.5 Mean Square Error	10
3 Algorithms	11
3.1 Necessary and Sufficient Conditions	12
3.2 Placement of Sensors for Problem 1	15
3.3 Finding an Approximate Optimal Trajectory for Problem 2	16
3.4 Finding an Approximate Optimal Trajectory for Problem 3	19
4 Results	22
5 Conclusion	26
Bibliography	27

List of Figures

1.1	54 temperature sensors deployed in the Intel Berkeley Research lab between February 28 th and April 5 th , 2004 [1].	2
1.2	A single quadcopter can fly over a farm and measure the height of the crop using a LIDAR sensor.	3
3.1	Collecting n_α measurements at O suffices to make accurate predictions at all points inside disk D_2 (Sufficient condition). No number of measurements at O can ensure predictive accuracy on points outside disk D_1 (Necessary condition).	14
3.2	We cover the disk of radius $3r_{max}$ using disks of radii $\frac{1}{\alpha}r_{max}$ in lawn mover pattern which requires no more than double the minimum number of $\frac{1}{\alpha}r_{max}$ disks required to cover a $3r_{max}$ disk. $18\alpha^2$ disks suffice to cover the bigger disk whereas the minimum number of $\frac{1}{\alpha}r_{max}$ disks required is $\frac{\pi(3r_{max})^2}{\pi(\frac{1}{\alpha}r_{max})^2} = 9\alpha^2$. The locations of disks of radii $\frac{1}{\alpha}r_{max}$ inside a disk of radius $3r_{max}$ are obtained by covering the square circumscribing bigger disk with smaller squares inscribed in smaller disks. The centers of smaller squares coincide with the centers of smaller disks.	17
3.3	Splitting the tour for one robot (τ) into 5 subtours. All robots start from the same initial position v_1	21
4.1	Actual and predicted OM content comparison. The farm is shown as the colored region with the colorbar denoting concentrations at different locations. All distance units are in <i>meter</i>	23
4.2	The mean percentage difference between the empirical MSE and the posterior variance. The mean was calculated over a set of grid locations inside the farm.	24
4.3	Disk placement covering the farm calculated by <i>DC</i> algorithm. The disks are of radii $3r_{max}$ which are concentric with disks of radii r_{max} in \mathcal{I} . The approximate optimal TSP tour visiting the centers is shown in blue. The lawn-mower detours have been omitted to make the figure more legible.	24
4.4	DiskCoverTour performs comparably with entropy-based and MI-based strategies. The shaded regions correspond to the standard deviation taken over five trials.	25
4.5	All robots start from and return to the depot after making measurements.	25

List of Abbreviations

k-DCT *k*-DISKCOVERTOUR

DC DISKCOVER

DCT DISKCOVERTOUR

GP Gaussian Process

IPP Informative Path Planning

Chapter 1

Introduction

Sensing, modeling, and tracking various spatially varying entities can improve our knowledge and understanding of them. This can have significant economic, environmental, and health implications. For example, knowing the content of various nutrients in the soil within a farm can help the farmers better understand soil chemistry. Understanding soil chemistry is helpful for the farmers to improve the yield and reduce the application of fertilizers [2]. An outbreak of certain chemicals inside a water body may have a big impact on marine life. Knowing the spatial extent of the outbreak is necessary for effective control and mitigation strategies [3]. Understanding the spatial variation of rock minerals can help in efficient mining strategies [4]. In all such applications, a key first step is the collection of data using appropriate sensors which can then be used to build models of the underlying phenomenon. However, collecting data can be tedious and often requires careful human planning. Manual data collection can also be dangerous. Consider a volcano monitoring activity. Volcano monitoring helps to see where previous lava flows have gone and to see where previous ash fall has occurred. Volcanic ash is usually pulverized rocks and glass particles and potentially catastrophic for the people engaged in monitoring [5]. To alleviate the risks associated with manual data collection, we need an alternative. One alternative is the use of robots equipped with appropriate sensors to collect data.

There are many factors to take into account when deploying robots for data collection. Usually, a trade-off has to be made between the quantity of sensing resources (e.g., number of deployed robots, energy consumption, mission time) and the quality of data collected. The robots can be deployed to act as stationary (Figure 1.1) or mobile (Figure 1.2) sensors depending on the application. Deploying robots to function as mobile sensors is especially challenging because of the need for planning robot paths as well. While deploying a number of mobile robotic sensors, one needs to plan the most informative resource-constrained observation paths to minimize the uncertainty in modeling and tracking.

Planning informative resource-constrained observation paths for robot sensors to estimate a spatially varying entity, often known as Informative Path Planning (IPP) has received a lot of attention in the robotics community [6, 7, 8, 9, 10, 11]. IPP deals with the problem of deciding an autonomous robot path along which maximum possible information about a quantity of interest can be extracted while operating under a set of resource constraints. In this thesis, our quantity of interest is a spatially varying phenomenon, often represented by a spatial field¹. Generally, the underlying spatial field is specified by a probabilistic model.

¹In this thesis, a spatial field is a function, $f(x), x \in U$, that is defined over a spatial domain, $U \subset \mathbb{R}^2$.




Figure 1.1: 54 temperature sensors deployed in the Intel Berkeley Research lab between February 28th and April 5th, 2004 [1].

One of the commonly used probabilistic models is Gaussian Process (GP) [12].

A GP is a stochastic process such that any finite collection of random variables has a multivariate Gaussian distribution. Hence, GPs can be treated as an extension of multivariate Gaussian distributions to infinite-sized collections of real-valued variables. This extension allows thinking of GPs as distributions not just over random vectors but distributions over random functions. A function drawn from a GP is then essentially an extremely high-dimensional vector drawn from an extremely high-dimensional multivariate Gaussian. GPs have been shown to accurately model many spatial fields [13].

After modeling the underlying spatial field, the next task is to plan a robot path based on some metric of the model of that spatial field. The robot travels along a path dictated in this step. Several metrics can be used to perform this planning step. An information-theoretic metric such as mutual information, entropy, or variance is typically used as a criterion to drive the robot to sampling locations [14]. Generally, the information-theoretic metrics are submodular and hence, an approximation guarantee can be given on the performance of the resulting algorithms [6]. Submodularity has a close connection with the property of diminishing returns. The intuition is that if a metric or utility function is submodular, then the marginal utility from adding a new element to a set of goods is larger than the one obtained when the new element is added to a larger set. Unfortunately, the information-theoretic metrics such as entropy, mutual information, etc. are indirect and do not consider the accuracy of the predictions. Unlike these works, we study how to ensure that the GP predicted mean² is accurate and present a constant-factor approximation algorithm if the hyperparameters of the GP kernel do not change.

We use prediction error as the metric to perform the planning step. The quality of predictions is usually characterized by the prediction error [15]. A commonly used prediction

²We use *predicted* mean and *estimated* mean interchangeably since the function is independent of time.



Figure 1.2: A single quadcopter can fly over a farm and measure the height of the crop using a LIDAR sensor.

error criterion is the mean square error (MSE). Our goal in this work is to plan the informative paths such that the MSE at all locations is below a certain predefined threshold, Δ after collecting measurements using mobile sensors.

We study three related problems, that of finding measurement locations to make measurements, planning a tour for a single robot to visit those measurement locations, and planning tours for multiple mobile robots to ensure that the MSE is below Δ . The objective is to minimize the number of measurement locations in the first problem and the total tour time in the second problem. In the case of multiple robots, the objective is to minimize the maximum time taken among all the robots. The total tour time is given by the measurement time and the travel time between measurement locations. The measurement time depends on the number of measurements taken at each location as well as the time to take a single measurement. Depending on the sensor, the measurement time can be zero (e.g., cameras) or non-zero (e.g., soil probes measuring organic content). We show that a non-adaptive algorithm suffices to solve the problem and yields a polynomial-time constant-factor approximation to the optimal algorithm. While other algorithms have been proposed in the past for estimating spatial fields, this is the first result that provides the theoretical guarantees on the total time for ensuring predictive accuracy at all points. In this thesis, our main contributions include:

- introducing stationary sensor placement and mobile sensor algorithms for ensuring that the MSE at each location in the environment is below a predefined threshold,
- providing polynomial-time constant-factor approximation guarantees on their performance, and
- showing their performance on a real-world dataset comprising of organic matter con-

centrations at various locations within a farm.

Similar problems have been studied in the literature. For example, Yfantis et. al. [16] studied a stationary sensor problem but their approach is restricted to a limited number of designs and not general enough. Further, we are not aware of any existing theoretical guarantees on the mobile sensor problems studied in this thesis.

The rest of the thesis is organized as follows: In Chapter 2, we present a discussion on related works and background on the problems studied. In Chapter 3, we formally present the problems and their solutions. Simulation results are presented in the Chapter 4 followed by the conclusion and scope for the future work in Chapter 5.

A preliminary version of this work was presented at the 13th International Workshop on the Algorithmic Foundations of Robotics (WAFR'18) [17]. In the preliminary version we provided guarantees for the chance constraints of incorrect predictions. In current work, a more pragmatic performance criterion, MSE, is considered which leads to stronger theoretical guarantees. Also, an extension of the algorithms for the multi-robot case is presented.

Chapter 2

Related Work and Background

We begin by reviewing the related work in sensor placement where the goal is to cover a given environment using sensors placed at fixed locations and mobile sensing where sensors are allowed to move and collect measurements from different locations.

2.1 Stationary Sensor Placement

When monitoring a spatial phenomenon, such as temperature or humidity in an environment, selection of a limited number of sensors and their locations is an important problem. The goal in this problem is to select the best k out of n possible sensor locations and use the measurements from these to make predictions about the spatial phenomenon. The typical formulation of a sensor selection problem makes it NP-hard [18]. Previous work used global optimization techniques such as branch and bound to exactly solve this problem [19, 20]. However, these exact approaches are often computationally intensive.

One can solve the task as an instance of the art-gallery problem [21, 22]: Find the minimum set of guards inside a polygonal workspace from which the entire workspace is visible. However, this version of the problem only covers vision-based sensors and does not take into account noisy measurements [6].

An alternative approach from spatial statistics is to learn a model of the phenomenon, typically as a GP [13, 23]. The learned GP model can then be used to predict the effect of placing sensors at particular locations and thus optimize their placement. For a given GP model, many criteria including information-theoretic ones have been proposed to evaluate the quality of placement. Shewry and Wynn introduced the maximum entropy criterion [24] where the sensors are placed sequentially at the locations of highest entropy. Ko et al. [25] proposed a greedy algorithm by formulating the entropy maximization as maximizing the determinant of the covariance matrix. However, entropy criterion tends to place the sensors at the boundary of the environment thus wasting sensed information [26]. Mutual information (MI) can be used as well [23, 27, 28]. Krause et al. [6] study the problem of maximizing MI for optimizing sensor placement problem. They present a polynomial-time approximation algorithm with constant factor guarantee by exploiting submodularity [29]. Eventually, they show that MI criterion leads to improved accuracy with a fewer number of sensors compared to other common design criteria such as entropy [24], A-optimal, D-optimal, and E-optimal design [30].

The above-mentioned methods estimate the prediction error indirectly. Nguyen et al. [31] consider choosing a set of n potential sensor measurements such that the root mean square prediction error is minimized. They present an annealing based algorithm for the sensor selection problem. Their algorithm starts by selecting a potential subset of cardinality k from the entire population of sensor locations. After that, it iteratively attempts to substitute the members of the selected subset by its neighbors according to an optimization criterion.

Unfortunately, any of the criteria discussed above cannot directly make any guarantees on the MSE in predictions at each point in the environment. Instead, we design a sensor placement algorithm which results in an accurate reconstruction of the spatial field using the collected sensor measurements. Most works in the past have focused on optimizing an objective function (entropy, MI, etc.) given the resource constraints (limited energy, number of sensors, and time, etc.). We optimize the resource requirement given the objective constraint (MSE below a predefined threshold Δ), predictive accuracy more than a predefined threshold in our case.

2.2 Mobile Sensing

The goal in the mobile sensing problem, also known as Informative Path Planning (IPP), is to compute paths for robots acting as mobile sensors. Paths are being computed in order to accurately estimate some underlying phenomenon, typically a spatial field [32, 33]. A central problem in IPP is to identify the hotspots in a large-scale spatial field. Hotspots are the regions in which the spatial field measurements exceed a predefined threshold. In many applications, it is necessary to assess the spatial extent and shape of the hotspot regions accurately. Low et al. presented a decentralized active robotic exploration strategy for probabilistic classification/labeling of hotspots in a GP-based spatial field [34]. The time needed by their strategy is independent of the map resolution and the number of robots, thus making it practical for *in situ*, real-time active sampling. Another formulation in hotspot identification is that of level set identification [35].

Previous works on level set boundary estimation and tracking [36, 37, 38] have primarily focused on communication of the sensor nodes, without giving much attention to individual sampling locations. Bryan et al. [39] proposed the straddle heuristic, which selects sampling locations by trading off uncertainty and proximity to the desired threshold level, both estimated using GPs. However, no theoretical justification had been given for its use and its extension to composite functions [40]. Gotovos et al. proposed a level set estimation algorithm, which utilizes GPs to model the target function and exploits its inferred confidence bounds to drive the selection process. They provided an information-theoretic bound on the number of measurements needed to achieve a certain accuracy, when the underlying function is sampled from a GP [41].

In many mobile sensing problems, it is not enough to identify only a few specific regions but estimate the entire spatial field accurately. It can be formulated as a path planning prob-

lem to observe a spatial field at a set of sampling locations, and then making inference about the unobserved locations [42]. Choosing and visiting the sample locations so that one can have an accurate prediction (point prediction and/or prediction interval) is of great importance in soil science, agriculture, and air pollution monitoring [28]. The objective functions used are usually submodular and thus exhibit a diminishing returns property. Submodularity arises due to the fact that near by measurement locations are correlated [43]. Chekuri and Pal introduced a quasi-polynomial time algorithm [44] for maximizing a submodular objective along the path using a *recursive greedy strategy*. This algorithm was further extended by Binney et al. [45] for spatiotemporal fields using average variance reduction [46] as the objective function. Meliou et al. [47] presented a nonmyopic informative path planning to model spatiotemporal fields. Their adaptive discretization technique allows to trade off solution quality and computational cost, empirically providing good solutions in lesser time.

Zhang and Sukhatme proposed an adaptive sampling algorithm consisting of a set of static nodes and a mobile robot tasked to reconstruct a scalar field [48]. They assume that the mobile robot can communicate with all the static nodes and acquire sensor readings from them. Based on this information, a path planner generates a path such that the resulting integrated mean square error is minimized subject to the constraint that the boat has a finite amount of energy.

An important issue in designing robot paths is deciding the next measurement location [7, 49, 50, 51], often referred to as the exploration strategy. Traditionally, conventional sampling methods [52] such as raster scanning, simple random sampling, and stratified random sampling have been used for single-robot exploration. Low et al. presented an adaptive exploration strategy called adaptive cluster sampling. It was demonstrated to reduce mission time and yield more information about the environment [53]. Their strategy performs better than a baseline sampling scheme called systematic sampling [54] using root mean squared error as a metric. A different adaptive multi-robot exploration strategy called MASP was presented in [55] which performs both wide-area coverage and hotspot sampling using nonmyopic path planning. MASP allows for varying adaptivity and its performance is theoretically analyzed. Further, it was demonstrated to sample efficiently from a GP and *logGP*. However, the time complexity of implementing it depends on the map resolution, which limits its large-scale use. To alleviate this computational difficulty, an information-theoretic approach was presented in [56]. The time complexity of the new approach was independent of the map resolution and less sensitive to the increasing robot team size. Garnett et al. [57] considered the problem of active search, which is also about sequential sampling from a domain of two (or more) classes. Their goal was to sample as many points as possible from one of the classes.

Yilmaz et al. [58] solved the adaptive sampling problem using mixed integer linear programming. Popa et al. [50] posed the adaptive sampling problem as a sensor fusion problem within the extended Kalman filter framework. Hollinger and Sukhatme proposed a sampling-based motion planning algorithm that generates maximal informative trajectories for the mobile robots to observe their environment [9]. Their information gathering algorithm ex-

tends ideas from rapidly-exploring random graphs. Using branch and bound techniques, they achieve efficient optimization of information gathering while also allowing for operation in continuous space with motion constraints. Low et al. [59] presented two approaches to solve IPP for *in situ* active sensing of GP-based anisotropic spatial fields. Their proposed algorithms can trade-off active sensing performance with computational efficiency. Ling et al. [10] proposed a nonmyopic adaptive GP planning framework endowed with a general class of Lipschitz continuous reward functions. Their framework can unify some active learning/sensing and Bayesian optimization criteria and offer practitioners flexibility to specify choices for defining new tasks. Tan et al. [60] introduced the receding-horizon cross-entropy trajectory optimization. Their focus was to sample around regions that exhibit extreme sensory measurements and much higher spatial variability, denoted as the region of interest. They used GP-UCB [61] as the optimization criteria which helps in exploring initially and converging on regions of interest eventually.

A naive implementation of GP prediction scales poorly with increasing training dataset size. Sparse GP frameworks can overcome this problem by using only a subset of the data to provide accurate estimates. A state-of-the-art sparse GP variant is SPGP [62, 63, 64, 65]. The SPGP framework learns a pseudo subset that best summarizes the training data. Mishra et al. introduced an online IPP framework AdaPP [66] which uses SPGP.

2.3 Sensing with Multiple Robots

Mobile sensing can be made faster by distributing the task among several robots. Multi-robot systems are capable of doing complex tasks and have been widely used in environmental sampling [67], coverage [68]. Robots can use local communication or control laws to achieve some collective goals.

Singh et al. [7] proposed a sequential allocation strategy that uses GP regression, which can be used to extend any single robot planning algorithm for the multi-robot problem. Their procedure approximately generalizes any guarantees for the single-robot problem to the multi-robot case. However, the approach works only when MI is the optimization objective. Cao et al. [59] presented two approaches along with their complexity analysis addressing a trade-off between active sensing performance and time efficiency. Luo et al. [69] combined adaptive sampling with information-theoretic criterion into the coverage control framework for model learning and simultaneous locational optimization. They presented an algorithm allowing for collaboratively learning the generalized model of density function using a mixture of GPs with hyperparameters learned locally from each robot. Kemna et al. [70] created a decentralized coordination approach which first splits the environments into Voronoi partitions and makes each vehicle then run within their own partition. Other multi-robot approaches used in other domains, e.g. exploration and estimation with ground vehicles, include auction-based methods [71, 72, 73] and spatial segregation, typically through Voronoi partitioning [74, 75].

Tokekar et al. [11] presented a constant factor approximation algorithm for the case of accurately classifying each point in a spatial field. The first step in the algorithm is to determine potentially misclassified points and then to find a tour visiting neighborhoods of each potentially misclassified point. In this paper, we study a regression version of the problem where every point is of interest. We exploit the properties of GP and squared-exponential kernel to find a constant-factor approximation algorithm. Before the details of the algorithms, we review some relevant background and useful properties of GPs and MSE.

2.4 Gaussian Processes

GPs provide a way to perform Bayesian non-linear regression. Though GPs have been frequently used in the classification problems [11, 12], our focus will be restricted to their application as a regression tool. In GPs, a prior distribution $p(f)$ is defined directly over the function mappings. After observing the data \mathcal{D} , one calculates the posterior distribution $p(f|\mathcal{D})$ using the Bayes theorem over those function mappings.

In GP regression, the posterior variance at any test location x is given by,

$$\hat{\sigma}_{x|X}^2 = k(x, x) - \mathbf{k}(x, X) [\mathbf{K}(X, X) + \omega^2 \mathbf{I}]^{-1} \mathbf{k}(X, x), \quad (2.1)$$

where k is the covariance (kernel) function with $\mathbf{K}_{pq} = k(x_p, x_q)$ and ω^2 is the variance of additive i.i.d. Gaussian measurement noise [12]. Note that the posterior variance at a particular location x conditioned on set of observations at locations $X = \{x_1, \dots, x_n\}$ does not depend on the actual observation but only on the locations from where the observations are collected. Multiple observations at a particular location is equivalent to that location being counted as many times as the number of measurements at that location. The kernel is a function that measures the *similarity* between two measurement locations [12]. For our analysis and experiments, we use a squared-exponential kernel with signal variance, and length scale equal to σ_0^2 and l , respectively.

Since the posterior variance is a function of only the measurement locations, the posterior variance for all points in the environment can be computed *a priori*, if the measurement locations are known, even without making any observations. In many implementations [14, 33, 76], the hyperparameters for the kernel k are tuned online as more data is gathered. As such, the hyperparameters may change with the observed data and the posterior variance will depend on the data observed, which may require adaptive planning. We assume that the hyperparameters are estimated *a priori*. Nevertheless, one can perform sensitivity analysis of the presented algorithms by varying the hyperparameters [77, 78].

The posterior mean $\hat{\mu}_{x|X}$ at a location x is given by a weighted linear combination of the observed data,

$$\hat{\mu}_{x|X} = \mathbf{k}(x, X) [\mathbf{K}(X, X) + \omega^2 \mathbf{I}]^{-1} \mathbf{y}, \quad (2.2)$$

where $\mathbf{y} = \{y_1, \dots, y_n\}$ denotes the observations at locations $X = \{x_1, \dots, x_n\}$.

2.5 Mean Square Error

MSE measures the expected squared difference between an estimator and the parameter the estimator is designed to estimate [79]. The MSE at a location x for an estimator \hat{f} is,

$$MSE(\hat{f}(x)) = Var(\hat{f}(x)) + (E(\hat{f}(x)) - f(x))^2, \quad (2.3)$$

where the Equation 2.3 is the bias-variance expression for an estimator. GP predicted value $\hat{f}(x)$ at a location x is an unbiased estimator of the true value $f(x)$ [16] and has a normal distribution with mean given by Equation 2.2, and variance given by Equation 2.1. Among all the estimators, GP predicted value has the minimum MSE [80] and because they are unbiased, the MSE at a location x is equal to the posterior variance of the predicted value, *i.e.*,

$$MSE(x) = \hat{\sigma}_{x|X}^2. \quad (2.4)$$

where the bias term in Equation 2.3 vanishes since GPs are unbiased estimators. From Equation 2.4, one can deduce that MSE for GPs is essentially the posterior variance and hence, any guarantees for the posterior variance hold for MSE as well.

Chapter 3

Algorithms

In this chapter, we formally define the problems and the algorithms. We assume that the environment is a two dimensional area $U \subset \mathbb{R}^2$ and the underlying spatial field is an instance of a GP, F [16]. F has a zero mean function, *i.e.*, $E[F(x)] = 0, \forall x \in U$ and an isotropic covariance function of the form,

$$C_Z(x, x') = \sigma_0^2 \exp\left(-\frac{(x - x')^2}{2l^2}\right); \forall x, x' \in U, \quad (3.1)$$

defined by a squared-exponential kernel where the hyperparameters σ_0^2 and l are known *a priori*. Let X denote the set of measurement locations within U produced by an algorithm. In the placement problem, the goal is to minimize the cardinality of X whereas in the mobile version the goal is to minimize the time required to visit and obtain measurements at X . MSE measures the expected squared difference, *i.e.*, $E\left(\hat{f}(x) - f(x)\right)^2$ between an estimator and the parameter the estimator is designed to estimate [79].

Problem 1 (Placement). Find the minimum number of measurement locations in U , such that the MSE at each location in U is below Δ , *i.e.*,

$$\begin{aligned} &\text{minimize} && \text{number of measurement locations } |X| \\ &\text{subject to} && \text{MSE}(x) \leq \Delta, \forall x \in U \end{aligned}$$

where $\text{MSE}(x)$ is the mean square error value in prediction at a location $x \in U$.

In the mobile version, a robot may obtain more than one measurement at a given location. Therefore, the number of measurements can be more than $|X|$. Let $n(X)$ be the total number of measurements obtained.

Problem 2 (Mobile). Find the minimum time trajectory for a mobile robot that obtains a finite set of measurements at one or more locations in U , such that the MSE at each location in U less than Δ , *i.e.*,

$$\begin{aligned} &\text{minimize} && \text{len}(\tau) + \eta n(X), \\ &\text{subject to} && \text{MSE}(x) \leq \Delta, \forall x \in U. \end{aligned}$$

τ denotes the tour of the robot. Robot travels at unit speed, obtains one measurement in η units of time and obtains $n(X)$ total measurements.

When there are more than one robots, their tours can start at the same starting location (often referred to as a *depot*) or can start at different locations. In this paper we focus on the former case, since we want to find tours for a single deployment. The latter case is more appropriate when the robots must persistently monitor the environment.

Problem 3 (multi-robot). For k robots starting from a given starting location (depot), design a set of trajectories that collectively obtain a finite set of measurements at one or more locations in U , such that the MSE at each location in U is less than Δ , *i.e.*,

$$\begin{aligned} & \text{minimize} && \max_{i \in \{1, \dots, k\}} \text{len}(\tau_i) + \eta n(X_i), \\ & \text{subject to} && \text{MSE}(x) \leq \Delta, \forall x \in U. \end{aligned}$$

τ_i denotes the tour of the i^{th} robot and X_i the subset of measurement locations covered by the i^{th} robot. Robots travel at unit speed, obtain one measurement in η units of time. i^{th} robot obtains $n(X_i)$ total measurements.

The solution for Problem 1 is a subset of the solution to Problem 2. Further, the solution for Problem 3 is derived from the solution for Problem 2. The three algorithms build on top of each other by: (1) finding a finite number of measurement locations for the robot; (2) finding a tour to visit all the measurement locations; and (3) splitting the tour from step 2 in multiple sub-tours for k robots. We exploit the properties of squared-exponential kernel to find the measurement locations. By knowing the value at a certain point within some tolerance, values at nearby points can be predicted albeit up to a larger tolerance. Our algorithms can be generalized to any type of measuring sensors as long as they satisfy the assumptions listed previously. The algorithms can also be generalized to higher dimensional environments, even though we illustrate using 2D examples.

3.1 Necessary and Sufficient Conditions

We start by deriving necessary conditions on how far a test location can be from the nearest measurement location. A test location corresponds to a point in the environment where we would like to make a prediction. Next, we prove a sufficient condition that if every point in the environment where no measurement is obtained (*test* location) is sufficiently close to a measurement location, then we can make accurate predictions at each point.

Lemma 3.1 (Necessary Condition). *For any test location x , if the nearest measurement location is at a distance r_{max} away, and,*

$$r_{max} > \sqrt[4]{-\log\left(1 - \frac{\Delta}{\sigma_0^2}\right)}, \quad (3.2)$$

then it is not possible to bring down the MSE below Δ at x .

Proof. Consider the posterior variance $\hat{\sigma}_x^2(n)$ at x (which is also equal to the *MSE* at x from Equation 2.4) after collecting n measurements, possibly from different locations. A lower bound on $\hat{\sigma}_x^2(n)$ can be obtained by assuming that all measurements were being collected at the nearest location x_i to x . Assume that the nearest measurement location x_i is distance r away from x . Assuming all the n measurements collected at x_i , lower bound for posterior variance at x , $\hat{\sigma}_x^2$ can be calculated using Equation 2.1,

$$\hat{\sigma}_x^2(n) \geq \sigma_0^2 - [k(x, x_i), \dots, k(x, x_i)] \begin{bmatrix} \sigma_0^2 + \omega^2 & & \sigma_0^2 \\ & \ddots & \\ \sigma_0^2 & & \sigma_0^2 + \omega^2 \end{bmatrix}^{-1} \begin{bmatrix} k(x, x_i) \\ \vdots \\ k(x, x_i) \end{bmatrix}. \quad (3.3)$$

It is worth mentioning that the square matrix in Equation 3.3 is of order $n \times n$ since there are n measurements. Substituting the value for $k(x, x_i) = \sigma_0^2 \exp\left(-\frac{r^2}{2l^2}\right)$ in Equation 3.3 and performing the required matrix operations, we get,

$$\hat{\sigma}_x^2(n) \geq \sigma_0^2 \left(1 - \frac{\exp\left(-\frac{r_{max}^2}{l^2}\right)}{1 + \frac{\omega^2}{n\sigma_0^2}} \right). \quad (3.4)$$

Even if we had collected infinitely many measurements at the nearest location x_i , the posterior variance will still be lower bounded as,

$$\hat{\sigma}_x^2(n) > \lim_{n \rightarrow \infty} \sigma_0^2 \left(1 - \frac{\exp\left(-\frac{r_{max}^2}{l^2}\right)}{1 + \frac{\omega^2}{n\sigma_0^2}} \right), \quad (3.5)$$

$$> \sigma_0^2 \left(1 - \exp\left(-\frac{r_{max}^2}{l^2}\right) \right). \quad (3.6)$$

If the posterior variance at x even with the infinitely many measurement collected at the nearest measurement location x_i (Equation 3.6) is greater than Δ , *i.e.*,

$$\begin{aligned} \Delta < \sigma_0^2 \left(1 - \exp\left(-\frac{r_{max}^2}{l^2}\right) \right) &\implies \\ r_{max} > l \sqrt{-\log\left(1 - \frac{\Delta}{\sigma_0^2}\right)}, & \end{aligned} \quad (3.7)$$

then it is not possible to bring down the *MSE* at x below Δ in any circumstance. □

Lemma 3.2 (Sufficient Condition). *For a test location $x \in U$, if there exists a measurement location $x_i \in X$, r distance away from x with n measurements at x_i , such that,*

$$r \leq l \sqrt{-\log\left(\left(1 + \frac{\omega^2}{n\sigma_0^2}\right) \left(1 - \frac{\Delta}{\sigma_0^2}\right)\right)}, \quad (3.8)$$

*then GP predictions at x will be accurate, *i.e.*, *MSE* at x will be smaller than Δ .*




Figure 3.1: Collecting n_α measurements at O suffices to make accurate predictions at all points inside disk D_2 (Sufficient condition). No number of measurements at O can ensure predictive accuracy on points outside disk D_1 (Necessary condition).

Proof. From Equation 3.4, we have an expression for posterior variance at x . Taking the other measurement locations into consideration can not increase the posterior variance at x . Information never hurts [81]! To prove the sufficiency, we consider n measurements at x_i only and discard others knowing that the other locations can not increase the posterior variance at x . Bounding the expression in Equation 3.4 with Δ results in,

$$\Delta \geq \sigma_0^2 \left(1 - \frac{\exp\left(-\frac{r^2}{l^2}\right)}{1 + \frac{\omega^2}{n\sigma_0^2}} \right), \quad (3.9)$$

$$\exp\left(-\frac{r^2}{l^2}\right) \geq \left(1 + \frac{\omega^2}{n\sigma_0^2}\right) \left(1 - \frac{\Delta}{\sigma_0^2}\right), \quad (3.10)$$

$$r \leq l \sqrt{-\log\left(\left(1 + \frac{\omega^2}{n\sigma_0^2}\right) \left(1 - \frac{\Delta}{\sigma_0^2}\right)\right)}. \quad (3.11)$$

□

Lemma 3.2 gives a sufficient condition for GP predictions to be accurate at any given test location $x \in U$. The following lemma shows that a finite number of measurements n_α , are sufficient to ensure predictive accuracy in a smaller disk of radius $\frac{1}{\alpha}r_{max}$ around x_i , where $\alpha > 1$ (Figure 3.1).

Lemma 3.3. *Given a disk of radius $\frac{1}{\alpha}r_{max}$ centered at x_i , n_α measurements at x_i suffice to make accurate predictions for all points inside the disk, where,*

$$n_\alpha \geq \left\lceil \frac{\omega^2}{\sigma_0^2} \frac{1}{\left(1 - \frac{\Delta}{\sigma_0^2}\right)^{\frac{1}{\alpha^2}-1} - 1} \right\rceil. \quad (3.12)$$

Proof. We want a sufficiency condition on the number of measurements n_α inside a disk of radius $\frac{1}{\alpha}r_{max}$. Lemma 3.2 gives an upper bound on the radius of a disk such that all points inside the disk will be accurately predicted after n_α measurements at the center. We construct a disk (D_2 in Figure 3.1), whose radius is equal to $\frac{1}{\alpha}r_{max}$ such that,

$$\frac{1}{\alpha}r_{max} \leq l\sqrt{-\log\left(\left(1 + \frac{\omega^2}{n_\alpha\sigma_0^2}\right)\left(1 - \frac{\Delta}{\sigma_0^2}\right)\right)}. \quad (3.13)$$

Plugging in the value of r_{max} from Lemma 3.1, squaring both sides in Equation 3.13 and re-arranging for n_α gives the required bound stated in Lemma 3.3. Ceiling function in Equation 3.12 accounts for the fact that n_α is an integer. \square

A packing of disks of radius r_{max} gives a lower bound on the number of measurements required to ensure predictive accuracy. On the other hand, a covering of disks of radius $\frac{1}{\alpha}r_{max}$ gives us an upper bound on the number of measurements required. To solve Problem 1, what remains is to relate the upper and lower bound and present an algorithm to actually place the disks of radii $\frac{1}{\alpha}r_{max}$.

3.2 Placement of Sensors for Problem 1

We use an algorithm similar to the one presented by Tekdas and Isler [82]. The exact procedure is outlined in Algorithm 1.

Algorithm 1 DiskCover

- 1: **procedure**
 - 2: **Input:** An environment.
 - 3: **Output:** Measurement locations.
 - 4: **begin**
 1. Design a set \mathcal{X} of disks of radii r_{max} which covers the environment and calculate a Maximal Independent Set (MIS) \mathcal{I} of \mathcal{X} greedily *i.e.*, $\mathcal{I} = \text{MIS}(\mathcal{X})$.
 2. Place disks of radii $3r_{max}$ concentric with disks in \mathcal{I} . Let the set of $3r_{max}$ radii disks is $\bar{\mathcal{X}}$.
 3. Cover each disk in $\bar{\mathcal{X}}$ with disks of radii $\frac{1}{\alpha}r_{max}$ as shown in Figure 3.2 and label centers of all disks of radii $\frac{1}{\alpha}r_{max}$.
 4. Return all the labeled points in previous step as measurement locations.
 - 5: **end procedure**
-

Theorem 3.4. DiskCover (Algorithm 1) gives an $18\alpha^2$ -approximation for Problem 1 in polynomial time.

Proof. Denote the set of measurement locations computed by the optimal algorithm to solve the Problem 1 by X^* . The function MIS in Step 1 of Algorithm 1 computes a maximally independent set of disks: the disks in \mathcal{I} are mutually non-intersecting (independent) and every disk in $\mathcal{X} \setminus \mathcal{I}$ intersects with some disk in \mathcal{I} (maximal). The set \mathcal{I} can be computed by a simple polynomial greedy procedure: choose an arbitrary disk d from \mathcal{X} , add it to \mathcal{I} , remove all disks in \mathcal{X} which intersect d , and repeat the procedure until no such d exists.

An optimal algorithm will collect measurements from at least as many measurement locations as the cardinality of \mathcal{I} . This can be proved by contradiction. Suppose an algorithm visits measurement locations fewer than the number of disks in \mathcal{I} . In that case, there will exist at least one disk of radius r_{max} in \mathcal{I} which will not contain a measurement location. This means that there will be at least a point in that disk which will be more than r_{max} away from each measurement location. From Lemma 3.1, the robot can never make accurate predictions at that point and hence violating the constraint in Problem 1. Hence,

$$|\mathcal{I}| \leq |X^*|. \quad (3.14)$$

Every disk in \mathcal{X} intersects at least one disk in \mathcal{I} and hence, lies within $3r_{max}$ of the center of a disk in \mathcal{I} . As a result, $\bar{\mathcal{X}}$ disks cover all the \mathcal{X} disks and hence, the entire environment.

Collecting measurements from $18\alpha^2$ locations inside a $3r_{max}$ disk suffice to make accurate predictions in that disk (satisfying the Problem 1 constraint for points belonging to that disk) as illustrated in Figure 3.2. DiskCover collects measurement from $18\alpha^2$ such locations per disk in $\bar{\mathcal{X}}$. It collects measurements from a total of $18\alpha^2|\bar{\mathcal{X}}|$ locations, hence, satisfying the constraint for all points in the area covered by union of $\bar{\mathcal{X}}$ disks. Since, union of $\bar{\mathcal{X}}$ disks covers the entire environment, DiskCover satisfies the constraint for all points in the environment. Multiplying both sides of Equation 3.14 with $18\alpha^2$, we get, $18\alpha^2|\mathcal{I}| \leq 18\alpha^2|X^*|$. Note that $|\bar{\mathcal{X}}| = |\mathcal{I}|$. Hence,

$$18\alpha^2|\bar{\mathcal{X}}| \leq 18\alpha^2|X^*|, \quad (3.15)$$

$$n_{\text{DiskCover}} \leq 18\alpha^2|X^*|, \quad (3.16)$$

where, $n_{\text{DiskCover}}$ is the number of measurement locations for DiskCover. \square

3.3 Finding an Approximate Optimal Trajectory for Problem 2

The algorithm for Problem 2 builds on the algorithm presented in the previous section. The locations where measurements are to be made, become the locations that are to be visited by the robot. The robot must obtain at least n_α measurements at the center of each disk of radius $\frac{1}{\alpha}r_{max}$. A pseudo-code of the algorithm is presented in the Algorithm 2.

Theorem 3.5. *DiskCoverTour (Algorithm 2) yields a constant-factor approximation algorithm for Problem 2 in polynomial time.*

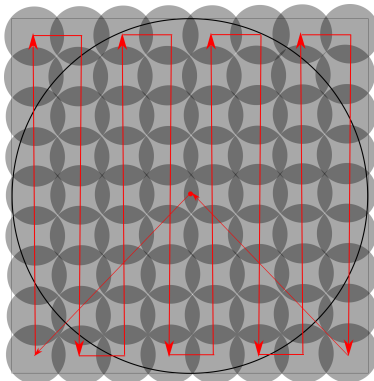


Figure 3.2: We cover the disk of radius $3r_{max}$ using disks of radii $\frac{1}{\alpha}r_{max}$ in lawn mover pattern which requires no more than double the minimum number of $\frac{1}{\alpha}r_{max}$ disks required to cover a $3r_{max}$ disk. $18\alpha^2$ disks suffice to cover the bigger disk whereas the minimum number of $\frac{1}{\alpha}r_{max}$ disks required is $\frac{\pi(3r_{max})^2}{\pi(\frac{1}{\alpha}r_{max})^2} = 9\alpha^2$. The locations of disks of radii $\frac{1}{\alpha}r_{max}$ inside a disk of radius $3r_{max}$ are obtained by covering the square circumscribing bigger disk with smaller squares inscribed in smaller disks. The centers of smaller squares coincide with the centers of smaller disks.

Algorithm 2 DiskCoverTour

1: **procedure**

2: **Input:** A set of measurement locations calculated from Algorithm 1.

3: **Output:** An approximate optimal tour visiting all the measurement locations.

4: **begin**

1. Calculate approximate TSP tour visiting centers of the $3r_{max}$ radius disks (set $\bar{\mathcal{X}}$) disks.
2. Cover $\bar{\mathcal{X}}$ disk containing the starting location in lawn-mover pattern visiting the centers of corresponding disks of radius $\frac{1}{\alpha}r_{max}$ and make n_α measurements at each center point.
3. Move to the center of next $\bar{\mathcal{X}}$ disk along the tour calculated in Step 1.
4. Repeat Steps 2 and 3 until all $\bar{\mathcal{X}}$ disks are covered.

5: **end procedure**

Proof. From Theorem 3.4, we have a constant approximation bound on number of measurement locations. Let the time (travel and measurement time) taken by the optimal algorithm be T^* . Using notation from Theorem 3.4, we assume that the optimal traveling salesperson with neighborhoods (TSPN) time to visit disks in \mathcal{I} be $T_{\mathcal{I}}^*$. In TSPN, we are given a set of geometric neighborhoods, and the objective is to find the shortest tour that visits at least one point in each neighborhood (disks in the present work) [11]. The optimal algorithm will visit at least all disks once in \mathcal{I} which gives the following minimum bounds on the optimal travel time (T_{travel}^*) and optimal measurement time ($T_{measure}^*$),

$$T_{\mathcal{I}}^* \leq T_{travel}^*; \eta|\mathcal{I}| \leq T_{measure}^*. \quad (3.17)$$

Let the optimal time to visit the centers of disks in \mathcal{I} be $T_{\mathcal{I}C}^*$. An upper bound on $T_{\mathcal{I}C}^*$ can be established by the fact that upon visiting each disk, the robot can visit the center of that disk and return back by adding an extra tour length of $2r_{max}$, *i.e.*, a detour of maximum length $|\mathcal{I}| \times 2r_{max}$ for all disks in \mathcal{I} . As a result: $T_{\mathcal{I}C}^* \leq T_{\mathcal{I}}^* + 2r_{max}|\mathcal{I}|$. Using inequality from Equation 3.17: $T_{\mathcal{I}C}^* \leq T_{travel}^* + 2r_{max}|\mathcal{I}|$. For any disk in $\bar{\mathcal{X}}$, the length of lawn-mover path starting from the its center and return back (Figure 3.2) after visiting all center points of $\frac{1}{\alpha}r_{max}$ disks will be of order $\mathcal{O}(\alpha^2)r_{max}$. Hence, the total travel time for DiskCoverTour is: $T_C + |\mathcal{I}|\mathcal{O}(\alpha^2)r_{max}$, where T_C is the $(1 + \epsilon)$ -approximated time with respect to the optimal TSP tour returned by the $(1 + \epsilon)$ -approximation algorithm to visit the centers of the disks in $\bar{\mathcal{X}}$ (or \mathcal{I} disks since they are concentric). T_C can be calculated in polynomial time [83] having bounds: $T_C \leq (1 + \epsilon) T_{\mathcal{I}C}^*$, with $T_{\mathcal{I}C}^*$ being the optimal TSP time to visit the centers of $\bar{\mathcal{X}}$ disks. Measurement time for DiskCoverTour is $18\alpha^2\eta n_2|\mathcal{I}|$. Hence, the total time T_{alg}^1 for DiskCoverTour is,

$$T_{alg}^1 = T_C + |\mathcal{I}|\mathcal{O}(\alpha^2)r_{max} + 18\alpha^2\eta n_2|\mathcal{I}| \quad (3.18)$$

$$\leq (1 + \epsilon) T_{\mathcal{I}C}^* + \mathcal{O}(\alpha^2)r_{max}|\mathcal{I}| + 18\alpha^2\eta n_2|\mathcal{I}|, \quad (3.19)$$

$$\leq (1 + \epsilon) (T_{travel}^* + 2r_{max}|\mathcal{I}|) + \mathcal{O}(\alpha^2)r_{max}|\mathcal{I}| + 18\alpha^2\eta n_2|\mathcal{I}|. \quad (3.20)$$

Length of any tour that visits k non-overlapping equal size disks of radii r is at least $0.24kr$ [84], which gives $0.24r_{max}|\mathcal{I}| \leq T_{\mathcal{I}}^*$. Combining this result with Equation 3.17 modifies the bounds in Equation 3.20 as,

$$T_{alg}^1 \leq \left((1 + \epsilon) \left(1 + \frac{2}{.24} \right) + \frac{\mathcal{O}(\alpha^2)}{.24} \right) T_{travel}^* + 18\alpha^2 n_2 T_{measure}^*, \quad (3.21)$$

$$\leq \max \left(9.33(1 + \epsilon) + \frac{82}{.24}, 72n_2 \right) (T_{travel}^* + T_{measure}^*), \quad (3.22)$$

$$\leq \max \left(9.33(1 + \epsilon) + \frac{\mathcal{O}(\alpha^2)}{.24}, 18\alpha^2 n_2 \right) T_1^* \leq cT_1^*, \quad (3.23)$$

where c denotes the corresponding maximum value in the right hand side of Equation 3.22, and is a constant. \square

3.4 Finding an Approximate Optimal Trajectory for Problem 3

When one robot can not handle a large territory, to speed up the task, k robots can be sent to collectively visit all measurement locations. A natural objective is to ensure that no robot has too large of a task. Hence, We choose our optimization criterion as minimizing the maximum of the k -robot tour costs. This is equivalent to minimizing the time taken by the last robot to return back to the common starting location. Our proposed algorithm only works if the robots start and return back to the same location – called depot. Any measurement location can be chosen as the depot but in our case, we assume that the robots start from and return back to a pre-defined depot.

We now describe an algorithm which employs a tour-splitting heuristic to plan for k robots. We modify the heuristic proposed by Frederickson et al. [85] to account for the measurement time, and not just the travel time.

Let the output tour of the robot from Algorithm 2 be denoted by τ and l_{max} be the distance of farthest measurement location from the depot.

Algorithm 3 k -DiskCoverTour

1: **procedure**

2: **Input:** Tour calculated from Algorithm 2, Depot location v_1 .

3: **Output:** k approximate optimal paths visiting all measurement locations collectively.

4: **begin**

1. For j^{th} robot, $1 \leq j < k$, find the last measurement location $v_{p(j)}$ such that the time taken to travel from v_1 to $v_{p(j)}$ along τ is not greater than $\frac{j}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2)$.
2. Obtain k subtours as $R_1 = (v_1, \dots, v_{p(1)}, v_1)$, $R_2 = (v_1, v_{p(1)+1}, \dots, v_{p(2)}, v_1)$, \dots
 $R_k = (v_1, v_{p(k-1)+1}, \dots, v_n, v_1)$.

5: **end procedure**

Theorem 3.6. k -DiskCoverTour (Algorithm 3) yields a $(c + 2)$ approximation algorithm for Problem 3 in polynomial time, given a c -approximation algorithm for Problem 2.

Proof. First, we prove that the time taken along every subtour is bounded and eventually show that the bound is within a constant factor of the optimal time. With k robots, let the subtours for 1^{st} and k^{th} robot are $v_1 \rightarrow v_{p(1)} \rightarrow v_1$ and $v_1 \rightarrow v_{p(k-1)+1} \rightarrow v_1$ respectively (an example with $k = 5$ is shown in Figure 3.3). Subtours for the remaining robots can be denoted by $v_1 \rightarrow v_{p(j-1)+1} \rightarrow v_{p(j)} \rightarrow v_1$, where $1 < j < k$.

Substituting $j = 1$ in Algorithm 3, the time to travel from v_1 to $v_{p(1)}$ along τ , $T(v_1 \xrightarrow{\tau} v_{p(1)})$ is no greater than $\frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2)$. Time for the first subtour is hence

bounded by $T(v_1 \xrightarrow{\tau} v_{p(1)}) + T(v_{p(1)} \rightarrow v_1)$, *i.e.*, $\frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2) + l_{max}$. From the condition in Algorithm 3, we know that $v_{p(k-1)}$ is the last location such that,

$$T\left(v_1 \xrightarrow{\tau} v_{p(k-1)}\right) \leq \frac{k-1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2), \quad (3.24)$$

and hence,

$$T\left(v_1 \xrightarrow{\tau} v_{p(k-1)+1}\right) \geq \frac{k-1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2). \quad (3.25)$$

Subtracting both sides from T_{alg}^1 ,

$$T_{alg}^1 - T\left(v_1 \xrightarrow{\tau} v_{p(k-1)+1}\right) \leq T_{alg}^1 - \frac{k-1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2), \quad (3.26)$$

which gives,

$$T\left(v_{p(k-1)+1} \xrightarrow{\tau} v_1\right) \leq \frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (3l_{max} + 2\eta n_2), \quad (3.27)$$

and hence, time for the last subtour is bounded by $T(v_1 \rightarrow v_{p(k-1)+1}) + T(v_{p(k-1)+1} \xrightarrow{\tau} v_1)$, *i.e.*, $\frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (3l_{max} + 2\eta n_2) + l_{max}$. Similar inequalities can be derived for remaining subtours as follows. For $1 \leq j \leq k-2$, following inequalities hold from Algorithm 3,

$$T(v_1 \xrightarrow{\tau} v_{p(j)+1}) \geq \frac{j}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2), \quad (3.28)$$

$$T(v_1 \xrightarrow{\tau} v_{p(j+1)}) \leq \frac{j+1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (l_{max} + \eta n_2), \quad (3.29)$$

Subtracting Equation 3.28 from 3.29 results in,

$$T(v_{p(j)+1} \xrightarrow{\tau} v_{p(j+1)}) \leq \frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)), \quad (3.30)$$

i.e., time taken along remaining subtours, $T(v_1 \rightarrow v_{p(j-1)+1} \xrightarrow{\tau} v_{p(j)} \rightarrow v_1)$, where $1 < j < k$, is also bounded by $\frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + 2l_{max}$. Hence, we can conclude that time taken along each subtour does not exceed $\frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (4l_{max} + 2\eta n_2)$.




Figure 3.3: Splitting the tour for one robot (τ) into 5 subtours. All robots start from the same initial position v_1 .

Let T_{alg}^k be the time taken for largest of the k subtours generated by the Algorithm 3, and T_k^* be the cost of the largest subtour in an optimal solution to Problem 3. We have,

$$T_{alg}^k \leq \frac{1}{k} (T_{alg}^1 - (2l_{max} + \eta n_2)) + (4l_{max} + 2\eta n_2) \quad (3.31)$$

$$\leq \frac{T_{alg}^1}{k} + (2l_{max} + \eta n_2) \left(2 - \frac{1}{k}\right). \quad (3.32)$$

From the triangle inequality, $T_k^* \geq \frac{1}{k} T_1^*$. It is natural to think that at least one robot will have to go to the farthest location from the depot and come back from there after collecting n_2 measurements which gives us a lower bound on the output of the optimal algorithm, *i.e.*, $2l_{max} + \eta n_2 \leq T_k^*$. Combining these results with Equation 3.23, we get,

$$T_{alg}^k \leq \frac{c}{k} T_1^* + T_k^* \left(2 - \frac{1}{k}\right) \quad (3.33)$$

$$\leq \left(c + 2 - \frac{1}{k}\right) T_k^* \quad (3.34)$$

$$\leq (c + 2) T_k^*. \quad (3.35)$$

□

Chapter 4


Results

We evaluate our algorithms with other baseline strategies through simulations based on a real-world dataset [86] collected from a farm (Figure 4.1(a)). The dataset consists of organic matter (OM) measurements manually collected from several hundred locations within the farm. The simulated sensor returns a noisy version of this ground truth measurement with an additive Gaussian noise of standard deviation, $\omega = 0.19$. The squared-exponential kernel has two hyperparameters: namely length scale (l) and signal variance (σ_0^2). The values of l and σ_0 were estimated to be 8.33 meter and 18.87 respectively by minimizing the negative log-marginal likelihood of the manually collected data. We assume that the estimated values are the true values of the kernel hyperparameters. The maximum and minimum values of the underlying field are 54.6 *ppm* and 25.4 *ppm* respectively (Figure 4.1(a)) as shown by the respective colorbars. The final predicted OM content after performing inference using the measurements obtained is shown in Figure 4.1(b). Figure 4.1(c) shows a plot of the prediction error. The regions where the OM content changes sharply tend to be more erroneously predicted as shown by the lighter colored regions in Figure 4.1(c). This can be attributed to the inherent smoothness assumptions of a squared-exponential kernel. The measurement locations computed by the DiskCoverTour are shown in Figure 4.1(d).

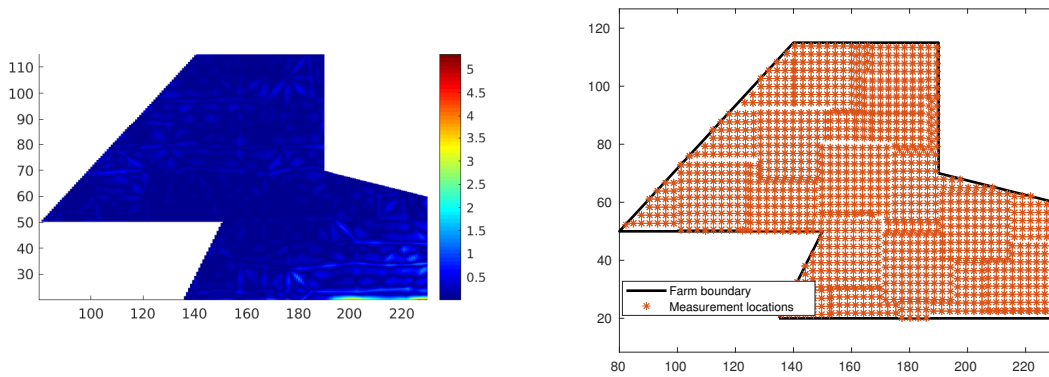
First, we verify our hypothesis that MSE is equal to the posterior variance for GPs. The MSE at a particular location is the expected squared error in prediction at that location. We calculate an empirical estimate of that expectation and anticipate that as the number of runs increases, the empirical estimate will become close to the actual MSE. A plot of the mean percent difference between the empirical MSE and the posterior variance is shown in Figure 4.2. As the number of trials increases, the mean difference between empirical MSE, which is essentially the MSE given enough number of trials, and the posterior variance decreases implying that the empirical MSE converges to the posterior variance asymptotically.

A covering of the farm with disks of radii $3r_{max}$ and an approximate optimal tour to visit the centers of those disks calculated by DiskCoverTour is shown in Figure 4.3. The approximate optimal TSP tour is calculated using $(1 + \epsilon)$ -approximation [83].

A comparison between DiskCoverTour, entropy-based, and MI-based strategies is shown in Figure 4.4. We observe that DiskCoverTour performs at par with the other strategies. The entropy-based strategy results in the most significant reduction in posterior variance and average empirical MSE initially. This can be explained by the fact that the entropy-based planning tends to spread the measurement locations far from each other resulting in



(a) The actual organic matter content (*ppm*). (b) The predicted organic matter content (*ppm*).



(c) Prediction error between the actual and the predicted OM content (*ppm*). (d) Measurement locations calculated by the DiskCover algorithm.

Figure 4.1: Actual and predicted OM content comparison. The farm is shown as the colored region with the colorbar denoting concentrations at different locations. All distance units are in *meter*.




Figure 4.2: The mean percentage difference between the empirical MSE and the posterior variance. The mean was calculated over a set of grid locations inside the farm.

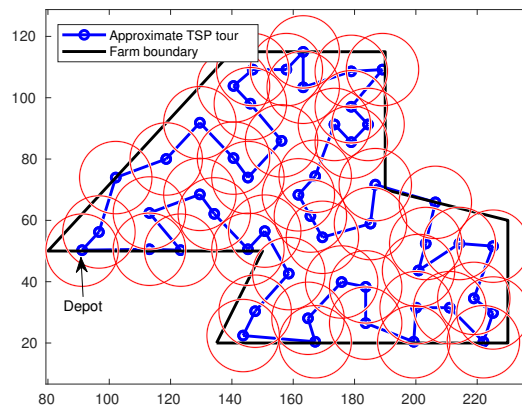


Figure 4.3: Disk placement covering the farm calculated by DC algorithm. The disks are of radii $3r_{max}$ which are concentric with disks of radii r_{max} in \mathcal{I} . The approximate optimal TSP tour visiting the centers is shown in blue. The lawn-mower detours have been omitted to make the figure more legible.

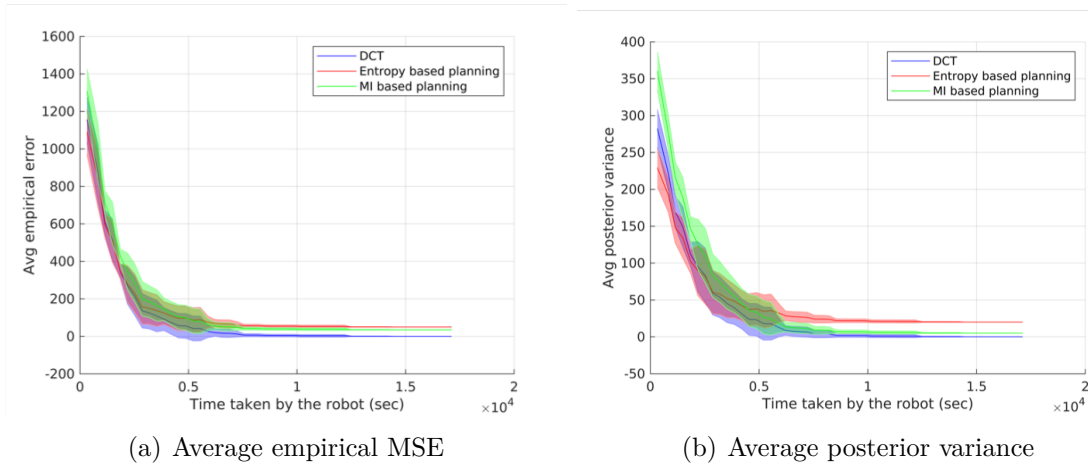


Figure 4.4: DiskCoverTour performs comparably with entropy-based and MI-based strategies. The shaded regions correspond to the standard deviation taken over five trials.

covering a bigger portion of the environment initially. However, DiskCoverTour converges to a lower value of average empirical MSE and average posterior variance.

For the multi-robot version, we assume that we have three robots. Splitting of a single robot tour (Figure 4.3) in three subtours is shown in Figure 4.5. The robots start from a common depot.

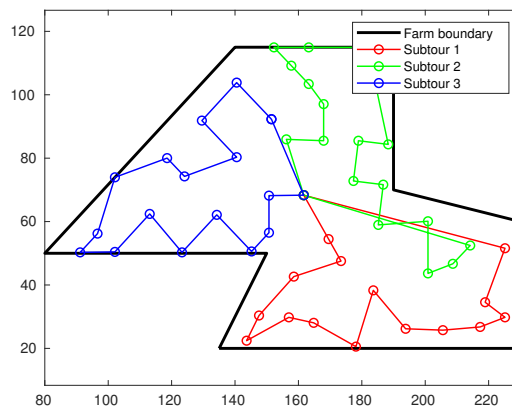


Figure 4.5: All robots start from and return to the depot after making measurements.

Chapter 5

Conclusion

In this thesis, we study several problems: Placing the minimum number of stationary sensors to track a spatial field, mapping a spatial field by a single as well as multiple robots while minimizing the time taken by the robots. For all the problems, lower bounds on the performance of any algorithm (including optimal) to solve respective problems are provided. We also provide polynomial-time approximation algorithms whose performances are within a constant factor of the lower bound and hence, within the respective optimal algorithms. We show that it is possible to learn a given spatial field accurately with high confidence without planning adaptively. Note that, if the kernel parameters are optimized online, then, one would require an adaptive strategy. Even though the algorithms suggested in this thesis perform comparatively with the strategies developed earlier, our algorithms have theoretical bounds on their performance. Our ongoing work is on developing competitive strategies for such cases as well as for spatio-temporal learning. Our ongoing work focuses on ensuring the present guarantees for adaptive cases and designing algorithms which can selectively make accurate predictions in the specific regions of interest (often called hot spots [60, 87]) inside a given environment.

Bibliography

- [1] Carlos Guestrin, Peter Bodik, Romain Thibaux, Mark Paskin, and Samuel Madden. Distributed regression: an efficient framework for modeling sensor network data. In *Proceedings of the 3rd international symposium on Information processing in sensor networks*, pages 1–10. ACM, 2004.
- [2] Wasim Aktar, Dwaipayan Sengupta, and Ashim Chowdhury. Impact of pesticides use in agriculture: their benefits and hazards. *Interdisciplinary toxicology*, 2(1):1–12, 2009.
- [3] L Adzibli and D Yuewen. Assessing the impact of oil spills on marine organisms. *J Oceanogr Mar Res*, 6(179):2, 2018.
- [4] Jan Blachowski. Spatial analysis of the mining and transport of rock minerals (aggregates) in the context of regional development. *Environmental Earth Sciences*, 71(3):1327–1338, Feb 2014.
- [5] G. Werner-Allen, K. Lorincz, M. Ruiz, O. Marcillo, J. Johnson, J. Lees, and M. Welsh. Deploying a wireless sensor network on an active volcano. *IEEE Internet Computing*, 10(2):18–25, March 2006.
- [6] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(Feb):235–284, 2008.
- [7] Amarjeet Singh, Andreas Krause, Carlos Guestrin, and William J Kaiser. Efficient informative sensing using multiple robots. *Journal of Artificial Intelligence Research*, 34:707–755, 2009.
- [8] Ruofei Ouyang, Kian Hsiang Low, Jie Chen, and Patrick Jaillet. Multi-robot active sensing of non-stationary gaussian process-based environmental phenomena. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 573–580. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [9] Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.
- [10] Chun Kai Ling, Kian Hsiang Low, and Patrick Jaillet. Gaussian process planning with lipschitz continuous reward functions: Towards unifying bayesian optimization, active learning, and beyond. In *AAAI*, pages 1860–1866, 2016.

- [11] Pratap Tokekar, Joshua Vander Hook, David Mulla, and Volkan Isler. Sensor planning for a symbiotic UAV and UGV system for precision agriculture. *IEEE Transactions on Robotics*, 32(6):1498–1511, 2016.
- [12] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [13] Noel Cressie. Statistics for spatial data. *Terra Nova*, 4(5):613–617, 1992.
- [14] Andreas Krause, Jure Leskovec, Carlos Guestrin, Jeanne VanBriesen, and Christos Faloutsos. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management*, 134(6):516–526, 2008.
- [15] Linh Van Nguyen, Sarath Kodagoda, Ravindra Ranasinghe, and Gamini Dissanayake. Simulated annealing based approach for near-optimal sensor selection in gaussian processes. In *2012 International Conference on Control, Automation and Information Sciences (ICCAIS)*, pages 142–147. IEEE, 2012.
- [16] Evangelos A Yfantis, George T Flatman, and Joseph V Behar. Efficiency of kriging estimation for square, triangular, and hexagonal grids. *Mathematical Geology*, 19(3):183–205, 1987.
- [17] Varun Suryan and Pratap Tokekar. Learning a spatial field with gaussian process regression in minimum time. In *The 13th International Workshop on the Algorithmic Foundations of Robotics*, 2018.
- [18] Fang Bian, David Kempe, and Ramesh Govindan. Utility based sensor selection. In *Proceedings of the 5th international conference on Information processing in sensor networks*, pages 11–18. ACM, 2006.
- [19] William J Welch. Branch-and-bound search for experimental designs based on d optimality and other criteria. *Technometrics*, 24(1):41–48, 1982.
- [20] Scott Kirkpatrick, C Daniel Gelatt, and Mario P Vecchi. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- [21] Dorit S Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *Journal of the ACM (JACM)*, 32(1):130–136, 1985.
- [22] H. González-Banos. A randomized art-gallery algorithm for sensor placement. In *Proceedings of the Seventeenth Annual Symposium on Computational Geometry, SCG '01*, pages 232–240, New York, NY, USA, 2001. ACM.
- [23] William F Caselton and James V Zidek. Optimal monitoring network designs. *Statistics & Probability Letters*, 2(4):223–227, 1984.

- [24] Michael C Shewry and Henry P Wynn. Maximum entropy sampling. *Journal of applied statistics*, 14(2):165–170, 1987.
- [25] Chun-Wa Ko, Jon Lee, and Maurice Queyranne. An exact algorithm for maximum entropy sampling. *Operations Research*, 43(4):684–691, 1995.
- [26] Naren Ramakrishnan, Chris Bailey-Kellogg, Satish Tadepalli, and Varun N Pandey. Gaussian processes for active data mining of spatial aggregates. In *Proceedings of the 2005 SIAM International Conference on Data Mining*, pages 427–438. SIAM, 2005.
- [27] Dale L Zimmerman. Optimal network design for spatial prediction, covariance parameter estimation, and empirical prediction. *Environmetrics*, 17(6):635–652, 2006.
- [28] Zhengyuan Zhu and Michael L. Stein. Spatial sampling design for prediction with estimated parameters. *Journal of Agricultural, Biological, and Environmental Statistics*, 11(1):24, Mar 2006.
- [29] George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*, 14(1):265–294, 1978.
- [30] Anthony C Atkinson. Optimal design. *Wiley StatsRef: Statistics Reference Online*, pages 1–17, 2014.
- [31] L. Van Nguyen, S. Kodagoda, R. Ranasinghe, and G. Dissanayake. Simulated annealing based approach for near-optimal sensor selection in gaussian processes. In *2012 International Conference on Control, Automation and Information Sciences (ICCAIS)*, pages 142–147, Nov 2012.
- [32] Syed Talha Jawaid and Stephen L Smith. Informative path planning as a maximum traveling salesman problem with submodular rewards. *Discrete Applied Mathematics*, 186:112–127, 2015.
- [33] Andreas Krause and Carlos Guestrin. Nonmyopic active learning of gaussian processes: an exploration-exploitation approach. In *Proceedings of the 24th international conference on Machine learning*, pages 449–456. ACM, 2007.
- [34] Kian Hsiang Low, Jie Chen, John M Dolan, Steve Chien, and David R Thompson. Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 105–112. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- [35] Frédéric Galland, Philippe Réfrégier, and Olivier Germain. Synthetic aperture radar oil spill segmentation by stochastic complexity minimization. *IEEE Geoscience and Remote Sensing Letters*, 1(4):295–299, 2004.

- [36] Aarti Singh, Robert Nowak, and Parmesh Ramanathan. Active learning for adaptive mobile sensing networks. In *Proceedings of the 5th international conference on Information processing in sensor networks*, pages 60–68. ACM, 2006.
- [37] Karthik Dantu and Gaurav S Sukhatme. Detecting and tracking level sets of scalar fields using a robotic sensor network. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 3665–3672. IEEE, 2007.
- [38] Sumana Srinivasan, Krithi Ramamritham, and Purushottam Kulkarni. Ace in the hole: Adaptive contour estimation using collaborating mobile sensors. In *Proceedings of the 7th international conference on Information processing in sensor networks*, pages 147–158. IEEE Computer Society, 2008.
- [39] Brent Bryan, Robert C Nichol, Christopher R Genovese, Jeff Schneider, Christopher J Miller, and Larry Wasserman. Active learning for identifying function threshold boundaries. In *Advances in neural information processing systems*, pages 163–170, 2006.
- [40] Brent Bryan and Jeff Schneider. Actively learning level-sets of composite functions. In *Proceedings of the 25th international conference on Machine learning*, pages 80–87. ACM, 2008.
- [41] Alkis Gotovos, Nathalie Casati, Gregory Hitz, and Andreas Krause. Active learning for level set estimation. In *Twenty-Third International Joint Conference on Artificial Intelligence*, 2013.
- [42] M. Wallat, T. Graepel, and K. Obermayer. Gaussian process regression: active data selection and test point rejection. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 3, pages 241–246 vol.3, July 2000.
- [43] Andreas Krause and Carlos Guestrin. Submodularity and its applications in optimized information gathering. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(4):32, 2011.
- [44] Chandra Chekuri and Martin Pal. A recursive greedy algorithm for walks in directed graphs. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 245–253. IEEE, 2005.
- [45] Jonathan Binney, Andreas Krause, and Gaurav S Sukhatme. Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8):873–888, 2013.
- [46] Abhimanyu Das and David Kempe. Algorithms for subset selection in linear regression. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 45–54. ACM, 2008.

- [47] Alexandra Meliou, Andreas Krause, Carlos Guestrin, and Joseph M. Hellerstein. Non-myopic informative path planning in spatio-temporal models. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1, AAAI'07*, pages 602–607. AAAI Press, 2007.
- [48] Bin Zhang and Gaurav S Sukhatme. Adaptive sampling for estimating a scalar field using a robotic boat and a sensor network. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 3673–3680. IEEE, 2007.
- [49] Naomi Ehrich Leonard, Derek A Paley, Francois Lekien, Rodolphe Sepulchre, David M Fratantoni, and Russ E Davis. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE*, 95(1):48–74, 2007.
- [50] Dan O Popa, Muhammad F Mysorewala, and Frank L Lewis. EKF-based adaptive sampling with mobile robotic sensor nodes. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2451–2456. IEEE, 2006.
- [51] Robert Sim and Nicholas Roy. Global a-optimal robot exploration in slam. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pages 661–666. IEEE, 2005.
- [52] Mohammad Rahimi, Richard Pon, William J Kaiser, Gaurav S Sukhatme, Deborah Estrin, and Mani Srivastava. Adaptive sampling for environmental robotics. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 4, pages 3537–3544. IEEE, 2004.
- [53] Kian Hsiang Low, Geoffrey J Gordon, John M Dolan, and Pradeep Khosla. Adaptive sampling for multi-robot wide-area exploration. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 755–760. IEEE, 2007.
- [54] A Tuchscherer. Thompson, steven k.: Sampling. john wiley & sons, inc., new york–chichester–brisbane–toronto–singapore 1992, 343pp.,£ 41, 95. *Biometrical Journal*, 35(7):848–848, 1993.
- [55] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Adaptive multi-robot wide-area exploration and mapping. In *Proceedings of the 7th International Joint Conference on Autonomous agents and Multiagent systems-Volume 1*, pages 23–30. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [56] Kian Hsiang Low, John M Dolan, and Pradeep Khosla. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Nineteenth International Conference on Automated Planning and Scheduling*, 2009.
- [57] Roman Garnett, Yamuna Krishnamurthy, Xuehan Xiong, Jeff Schneider, and Richard Mann. Bayesian optimal active search and surveying. *arXiv preprint arXiv:1206.6406*, 2012.

- [58] Namik Kemal Yilmaz, Constantinos Evangelinos, Pierre FJ Lermusiaux, and Nicholas M Patrikalakis. Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4):522–537, 2008.
- [59] Nannan Cao, Kian Hsiang Low, and John M Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, pages 7–14. International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [60] Yew Teck Tan, Abhinav Kunapareddy, and Marin Kobilarov. Gaussian process adaptive sampling using the cross-entropy method for environmental sensing and monitoring. In *International Conference on Robotics and Automation 2018*, pages 6220–6227, 05 2018.
- [61] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, May 2012.
- [62] Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in neural information processing systems*, pages 1257–1264, 2006.
- [63] Lehel Csató and Manfred Opper. Sparse on-line gaussian processes. *Neural computation*, 14(3):641–668, 2002.
- [64] Matthias Seeger, Christopher Williams, and Neil Lawrence. Fast forward selection to speed up sparse gaussian process regression. *Artificial Intelligence and Statistics 9*, 2003.
- [65] Trong Nghia Hoang, Quang Minh Hoang, and Kian Hsiang Low. A unifying framework of anytime sparse gaussian process regression models with stochastic variational inference for big data. In *Proceedings of the 32Nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, pages 569–578. JMLR.org, 2015.
- [66] Rajat Mishra, Mandar Chitre, and Sanjay Swarup. Online informative path planning using sparse gaussian processes. In *2018 OCEANS-MTS/IEEE Kobe Techno-Oceans (OTO)*, pages 1–5. IEEE, 2018.
- [67] Wenhao Luo, Shehzaman S Khatib, Sasanka Nagavalli, Nilanjan Chakraborty, and Katia Sycara. Distributed knowledge leader selection for multi-robot environmental sampling under bandwidth constraints. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5751–5757. IEEE, 2016.

- [68] Jorge Cortes, Sonia Martinez, Timur Karatas, and Francesco Bullo. Coverage control for mobile sensing networks. *IEEE Transactions on robotics and Automation*, 20(2):243–255, 2004.
- [69] Wenhao Luo and Katia Sycara. Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes. In *IEEE International Conference on Robotics and Automation (ICRA)*, May 2018.
- [70] Stephanie Kemna, John G Rogers, Carlos Nieto-Granda, Stuart Young, and Gaurav S Sukhatme. Multi-robot coordination through dynamic voronoi partitioning for informative adaptive sampling in communication-constrained environments. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2124–2130. IEEE, 2017.
- [71] Robert Zlot, Anthony Stentz, M Bernardine Dias, and Scott Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No. 02CH37292)*, volume 3, pages 3016–3023. IEEE, 2002.
- [72] Reid G. Simmons, David Apfelbaum, Wolfram Burgard, Dieter Fox, Mark Moors, Sebastian Thrun, and Håkan L. S. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 852–858. AAAI Press, 2000.
- [73] Weihua Sheng, Qingyan Yang, Song Ci, and Ning Xi. Multi-robot area exploration with limited-range communications. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 2, pages 1414–1419. IEEE, 2004.
- [74] Daniel E Soltero, Mac Schwager, and Daniela Rus. Generating informative paths for persistent sensing in unknown environments. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2172–2179. IEEE, 2012.
- [75] Alessandro Marino, Gianluca Antonelli, Antonio Pedro Aguiar, António Pascoal, and Stefano Chiaverini. A decentralized strategy for multirobot sampling/patrolling: Theory and experiments. *IEEE Transactions on Control Systems Technology*, 23(1):313–322, 2015.
- [76] Alexandra Meliou, Andreas Krause, Carlos Guestrin, and Joseph M. Hellerstein. Non-myopic informative path planning in spatio-temporal models. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1, AAAI’07*, pages 602–607. AAAI Press, 2007.

- [77] Marc C Kennedy and Anthony O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001.
- [78] Jeremy E Oakley and Anthony O’Hagan. Probabilistic sensitivity analysis of complex models: a bayesian approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):751–769, 2004.
- [79] Trevor Hastie, Robert Tibshirani, Jerome Friedman, and James Franklin. The elements of statistical learning: data mining, inference and prediction. *The Mathematical Intelligencer*, 27(2):83–85, 2005.
- [80] Johan Wågberg, Dave Zachariah, Thomas B Schön, and Petre Stoica. Prediction performance after learning in gaussian process regression. *arXiv preprint arXiv:1606.03865*, 2016.
- [81] Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
- [82] Onur Tekdas and Volkan Isler. Sensor placement for triangulation-based localization. *IEEE transactions on Automation Science and Engineering*, 7(3):681–685, 2010.
- [83] S. Arora. Nearly linear time approximation schemes for euclidean tsp and other geometric problems. In *Proceedings 38th Annual Symposium on Foundations of Computer Science*, pages 554–563, Oct 1997.
- [84] Onur Tekdas, Deepak Bhadauria, and Volkan Isler. Efficient data collection from wireless nodes under the two-ring communication model. *The International Journal of Robotics Research*, 31(6):774–784, 2012.
- [85] G. N. Frederickson, M. S. Hecht, and C. E. Kim. Approximation algorithms for some routing problems. In *17th Annual Symposium on Foundations of Computer Science (sfcs 1976)*, pages 216–227, Oct 1976.
- [86] D. J. Mulla, A. C. Sekely, and M. Beatty. Evaluation of remote sensing and targeted soil sampling for variable rate application of nitrogen. In P. C. Robert, R. H. Rust, and W. E. Larson, editors, *Proceedings of the 5th International Conference on Precision Agriculture*, pages 1–15, Madison, USA, 2000. American Society of Agronomy.
- [87] Andreas Krause, H Brendan McMahan, Carlos Guestrin, and Anupam Gupta. Robust submodular observation selection. *Journal of Machine Learning Research*, 9(Dec):2761–2801, 2008.