Learning a Spatial Field with Gaussian Process Regression in Minimum Time*

Varun Suryan and Pratap Tokekar

Department of Electrical and Computer Engineering, Virginia Tech, USA
{suryan, tokekar}@vt.edu

Abstract. We study an informative path planning problem where the goal is to minimize the time required to learn a spatial field using Gaussian Process (GP) regression. Specifically, given parameters $0 < \epsilon, \delta < 1$, our goal is to ensure that the predicted value at all points in an environment lies within $\pm\epsilon$ of the true value with probability at least δ. We study two versions of the problem. In the sensor placement version, the objective is to minimize the number of sensors placed. In the mobile sensing version, the objective is to minimize the total travel time required to visit the sensing locations. The total time is given by the time spent obtaining measurements as well as time to travel between measurement locations. By exploiting the smoothness properties of GP regression, we present constant-factor approximation algorithms for both problems that make accurate predictions at each point. Our algorithm is a deterministic, non-adaptive one and based on the Traveling Salesperson Problem. In addition to theoretical results, we also compare the empirical performance using a real-world dataset with other baseline strategies.

1 Introduction

The problem of planning adaptive, informative paths for mobile sensors to estimate a spatial field[1] has received a lot of attention in the community recently [1–6]. The underlying field is typically modeled as a Gaussian Process (GP) [7] and estimated through GP regression [2] or GP classification [6]. An information-theoretic measure such as mutual information, entropy, or variance is used as the criterion to drive the robot to sampling locations [8]. Optimizing such measures leads to adaptive sampling of the field. Typically an approximation guarantee on the information-theoretic measure can be given, especially if the measure used is a submodular one [1]. Such measures are indirect since they are based on the variance of the prediction but do not consider the accuracy of the predicted mean. Unlike these works, we study how to ensure that the predicted

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

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

mean[2] is accurate and present a constant-factor approximation algorithm if the hyper-parameters of the kernel do not change.

We use the notion of chance-constrained optimization [9] to find the optimal trajectory that ensures that the probability of incorrect estimation is less than a user-specified threshold. Let $f(x)$ be the true function value and $\hat{\mu}(x)$ be the estimated value at a point x. The estimated value is said to be chance-constrained, if $Pr\{|\hat{\mu}(x) - f(x)| \leq \epsilon\} > \delta$ for all x. Here, $0 < \delta, \epsilon < 1$ are the desired confidence (*i.e.,* "probably") and accuracy (*i.e.,* "approximately") parameters. Note that this is a stronger requirement in the sense that the prediction at each point must satisfy the condition for the algorithm. This is in contrast to aggregate measures such as entropy or mutual information that consider the sum or average measure over all the points. Furthermore, we specifically consider the mean predicted value rather than just minimizing the uncertainty in prediction.

We study two related problems, that of finding measurement locations to place sensors and planning a tour for a single mobile sensor to ensure chance-constrained GP regression. The objective functions are to minimize the number of measurement locations and the total tour time, respectively. The total 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 a non-adaptive algorithm suffices to solve the problem and yields a polynomial-time constant-factor approximation to the optimal. While other algorithms have been proposed in the past for estimating spatial fields, this is the first result that gives guarantees on the total travel time for ensuring predictive accuracy at all points.

The rest of the paper is organized as follow. We first present the work related to our problem (Section 2). Next, we formally define the actual problems we study in this paper (Section 3). The two algorithms, DiskCover and DiskCoverTour, along with their theoretical analysis are given in Section 4. We compare the performance of the algorithms with baseline strategies on a real-world dataset [6] in Section 5.

2 Related Work

We begin by reviewing the related work in the areas of 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

Sensor placement algorithms can be classified as *geometric* and *model-based* approaches [1] broadly, based upon the assumptions made about the environment.

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

In geometric approaches, typically the goal is to find a set of measurement locations which cover the environment [10, 11].

In model-based approaches, one takes a model of the environment (generally a probabilistic one) and finds measurement locations to optimize an objective function of that model. Information-theoretic measures, notably entropy and mutual information, are frequently used as the optimization criteria [12, 13]. A general approach is to greedily place sensors at locations where the uncertainty about the phenomena is highest. The entropy criterion tends to place sensors along the boundaries of the environment [14] which can be wasteful. An alternative criterion of mutual information seeks to find placements that are most informative about locations where no sensors are placed [12, 1]. This optimization criterion directly measures the effect of sensor placements on the posterior uncertainty of GP. In their seminal work, Krause et al. [1] studied the combinatorial optimization problem of placing sensors to maximize mutual information. The problem is NP-complete. They showed conditions under which mutual information is monotone and submodular [15] and exploited this result to prove that a greedy placement algorithm yields a constant-factor approximation in polynomial time. Unfortunately, entropy and mutual information cannot directly make any guarantees on the predictive accuracy at specific points in the environment. Instead, we design a sensor placement algorithm that guarantees predictive accuracy at every point in the environment.

2.2 Mobile Sensing

The problem of designing optimal trajectories and the related problem of selecting measurement locations has recently received considerable attention. Low et al. [16] presented a control law to minimize the probability of misclassification in a field estimated by GP regression. The authors enforced measurements to be taken continuously, and forced the sensors to only move along a four-connected grid. Song et al. [17] presented an algorithm to localize multiple radio sources using a mobile robot. They presented upper bounds on the time required to localize the sources up to a desired probability. Otte et al. [18] studied the problem of navigating to a global goal, while foraging for interesting locations along the way. They analyzed two greedy strategies to bound the total distance traveled and verified the bounds through simulations. Krause et al. proposed mutual information as a measure of uncertainty [1]. An algorithm to place sensing locations was given which can closely approximate the optimal increase in mutual information. The work was extended to mobile sensor routing in [19], and multiple robots in [2]. Tan et al. developed an adaptive sampling strategy using receding-horizon cross-entropy trajectory optimization on a GP [20].

Tokekar et al. [6] 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.

3 Problem Formulation

In this section, we formally define the terminology and the two problems studied. We assume that the environment is a compact domain $U \subset \mathbb{R}^2$. We make the following assumptions about the spatial function $f_x \equiv f(x)$:

Assumption 1 (Smoothness) *The true function is smooth in the sense of Lipschitz [21], i.e.,* $\forall x_i, x_j \in U : |f_{x_i} - f_{x_j}| \leq \mathcal{L}||x_i - x_j||_2$*, where* $\mathcal{L}$ *is the Lipschitz smoothness constant.*

Assumption 2 (Kernel) *The GP regression uses a squared-exponential kernel [7]. The hyperparameters of the kernel (length scale and signal variance) are known* a priori.

Assumption 3 (Bound on Measurements) *Optimal algorithms for both problems satisfy the chance constraints using a finite number of measurements, no more than* N*.*

Let X denote the set of measurement locations within U produced by an algorithm. In the placement problem, our goal is to minimize the cardinality of X whereas in the path planning problem the goal is to minimize the time required to visit and obtain measurements at X. The two problems are formally defined below.

Problem 1 (Placement). Given Assumptions 1–3, find the minimum number of measurement locations in U, such that the posterior GP prediction for any point in U is within $\pm\epsilon$ of the actual value with probability at least δ, *i.e.,*

$$\begin{aligned} &\text{minimize} && \text{number of measurement locations } |X| \\ &\text{subject to} && Pr\{|\hat{\mu}(x) - f(x)| \leq \epsilon\} \geq \delta, \forall x \in U \end{aligned}$$

where $f(x)$ is the actual function value that is to be estimated at a point $x \in U$, and $\hat{\mu}(x)$ is the predicted value using measurements obtained at locations X.

Problem 2 (Mobile). Given Assumptions 1–3, find a minimum time trajectory for a mobile sensor that obtains a finite set of measurements at one or more locations in U, such that the posterior GP prediction for any point in U is within $\pm\epsilon$ of the true value with probability at least δ, *i.e.,*

$$\begin{aligned} &\text{minimize} && len(\tau) + \eta\beta(X), \\ &\text{subject to} && Pr\{|\hat{\mu}(x) - f(x)| \leq \epsilon\} \geq \delta, \forall x \in U. \end{aligned}$$

τ denotes the tour of the robot. Assume that the robot travels at unit speed, obtains one measurement in η units of time and obtains $\beta(X)$ total measurements at the measurement locations X. Here, we use the function $\beta(X)$ to take into account the fact that the robot may obtain more than one measurement at a specific location. Therefore, the number of measurements can be more than $|X|$.

4 Algorithm

We present the details of our algorithm for both problems. The solution to Problem 1 is a subset of the solution to Problem 2.

Our algorithm consists of two main stages: (1) finding a finite number of measurement locations for the robot; and (2) finding a tour to visit all the sensing locations. We exploit the Lipschitz smoothness property to find the measurement locations. By knowing the value at a certain point within some tolerance, we can predict values at nearby points albeit up to a larger tolerance. Our algorithm can be generalized to any type of 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.

Before we present the detailed algorithms, we review some relevant background and useful properties on GPs.

4.1 Useful Properties of GPs

In GP regression, the posterior variance at a particular location x conditioned on set of observations at locations $X = \{x_1, \ldots, x_n\}$ does not depend on the data gathered but only on the locations where the data is gathered. The posterior variance on any location x is given by,

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

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 [7]. The kernel is a function that measures the *similarity* between two measurement locations [7]. 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 sensing 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 [8, 22, 19], 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 [23, 24].

The objectives for the two problems depend on not only the posterior variance but also the posterior mean. The posterior mean $\hat{\mu}_{x|X}$ at a location x is given by a weighted linear combination of the observed data plus a bias term,

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

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

4.2 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 1 (Necessary Condition). *For any test location x, if the nearest measurement location is at a distance r_{max} away, and,*

$$r_{max} \geq l\sqrt{-\log\left(\frac{\sigma_0^2+\omega^2}{N\sigma_0^2}\left(1-\frac{\epsilon}{\sqrt{2}\sigma_0^2\,\mathrm{erf}^{-1}(\delta)}\right)\right)}, \tag{3}$$

then it is not possible to make ϵ-accurate predictions at that test location with desired probability δ.

Proof. GP prediction at x is a normal random variable with mean $\hat{\mu}_{x|X}$ (Equation 2) and variance $\hat{\sigma}^2_{x|X}$ (Equation 1). Even if the posterior mean was exactly equal to the true value of the function, the posterior variance has to be small enough to ensure chance-constrained regression, *i.e.,* total probability mass has to be more than δ in $\pm\epsilon$-band around mean $\hat{\mu}_{x|X}$. A lower bound on $\hat{\sigma}^2_{x|X}$ can be obtained using the fact that independent measurements will result in highest decrease in uncertainty at x [25], *i.e.,*

$$\hat{\sigma}^2_{x|X} \geq \hat{\sigma}^2_{x|X_I}, \tag{4}$$

where, $\hat{\sigma}^2_{x|X_I}$ is the posterior variance at x assuming all the measurements were independent. Posterior variance at a test location x assuming all the n measurements independent is given by,

$$\hat{\sigma}^2_{x|X_I} = \sigma_0^2 - \begin{bmatrix}k(x,x_1),\ldots,k(x,x_n)\end{bmatrix}\begin{bmatrix}\sigma_0^2+\omega^2 & & \mathbf{0}\\ & \ddots & \\ \mathbf{0} & & \sigma_0^2+\omega^2\end{bmatrix}^{-1}\begin{bmatrix}k(x,x_1)\\ \vdots \\ k(x,x_n)\end{bmatrix}. \tag{5}$$

Non-diagonal entries of the co-variance matrix in Equation 5 are zero (independence among the measurement locations). Assuming the nearest measurement location to be r_{max} away from the test location x, we have,

$$k(x,x_i) \leq \sigma_0^2 \exp\left(-\frac{r^2_{max}}{2l^2}\right), \forall x_i \in X, \tag{6}$$

which combined with Equations 4, and 5 results in,

$$\hat{\sigma}^2_{x|X} \geq \hat{\sigma}^2_{x|X_I} \geq \sigma_0^2\left(1-\exp\left(-\frac{r^2_{max}}{l^2}\right)\frac{n\sigma_0^2}{\sigma_0^2+\omega^2}\right). \tag{7}$$

Now, by lower bounding the the probability mass, we have,

$$\int_{\hat{\mu}_{x|X}-\epsilon}^{\hat{\mu}_{x|X}+\epsilon} \mathcal{N}\left(\hat{\mu}_{x|X}, \hat{\sigma}^2_{x|X}\right) \geq \delta \implies \frac{\epsilon}{\sqrt{2}\,\mathrm{erf}^{-1}(\delta)} \geq \hat{\sigma}^2_{x|X}. \tag{8}$$

where erf^{-1} is the inverse error function [26]. Equation 8 gives us a necessary bound on the posterior variance at x. Combining Equation 8 with Equation 7 and using the fact, $N > n$, if r_{max} satisfies the inequality in Lemma 1, then it is not possible to satisfy Equation 8. Hence, it is impossible to ensure the chance-constraints for the GP regression. That is, if the nearest measurement location from a test location is more than r_{max} away then it is impossible to make ϵ-accurate predictions at the test location with probability δ. □

Lemma 2 (Sufficient Condition). *For a test location $x \in U$, if there exists a measurement location $x_i \in X$ with n measurements at x_i, such that,*

$$d(x, x_i) \leq \frac{1}{\mathcal{L} + \sqrt{2}\mathcal{L}_1}\left(\epsilon - \sqrt{\frac{2}{\frac{1}{\sigma_0^2} + \frac{n}{\omega^2}}}\,\mathrm{erf}^{-1}\left(\frac{\delta}{1 - a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right)}\right)\right), \tag{9}$$

then GP predictions at x will be ϵ-accurate with high confidence, i.e., $Pr\left[\left|\hat{f}_{x|X} - f_x\right| \leq \epsilon\right] \geq \delta$, *where $a, b,$* and *$\mathcal{L}_1$ are constants.*

Proof. Consider a measurement location, say x_i, that is a distance $d(x, x_i)$ away from the point x. From the Lipschitz smoothness assumption we have,

$$|f_x - f_{x_i}| \leq \mathcal{L} d(x, x_i). \tag{10}$$

A bound on GP posterior variance $\hat{\sigma}^2_{x_i|X}$ at x_i after making n measurements at x_i can be obtained using Equation 1,

$$\hat{\sigma}^2_{x_i|X} \leq \frac{\omega^2}{n + \frac{\omega^2}{\sigma_0^2}}. \tag{11}$$

The inequality in Equation 11 follows from the fact that right hand side does not consider the effect of measurement locations other than x_i which can not increase the posterior variance at x_i. GPs are consistent function estimators [7] and hence GP mean prediction at any measurement location converges to actual value of the function given sufficient number of measurements at that location. GP prediction $\hat{f}_{x_i|X}$ at x_i is a normal random variable with mean $\hat{\mu}_{x_i|X}$ and $\hat{\sigma}^2_{x_i|X}$, hence,

$$\frac{\hat{f}_{x_i|X} - \hat{\mu}_{x_i|X}}{\hat{\sigma}_{x_i|X}} \sim \mathcal{N}(0, 1). \tag{12}$$

Using consistency of GPs, Equation 12 can be written as,

$$\frac{\hat{f}_{x_i|X} - f_{x_i}}{\hat{\sigma}_{x_i|X}} \sim \mathcal{N}(0, 1) \implies Pr\left(\left|\hat{f}_{x_i|X} - f_{x_i}\right| \leq \frac{k\omega}{\sqrt{n + \frac{\omega^2}{\sigma_0^2}}}\right) = \mathrm{erf}\left(\frac{k}{\sqrt{2}}\right). \tag{13}$$

Squared-exponential kernel for one dimensional input satisfies the following high probability bound [21] on the derivatives of GP sample paths $\hat{f}_{x|X}$: for some constants $a, b > 0$,

$$Pr\left(\sup_{x\in U}\left|\partial\hat{f}_{x|X}/\partial x\right| < \mathcal{L}_1\right) \geq 1 - a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right). \tag{14}$$

Equation 14 bounds the first order derivative of the sampled functions with high probability. Hence, sampled functions are Lipschitz smooth with constant $\mathcal{L}_1$. As a result, GP predictions at any locations x and x_i satisfies,

$$Pr\left(\left|\hat{f}_x - \hat{f}_{x_i}\right| \leq \sqrt{2}\mathcal{L}_1 d(x, x_i)\right) \geq 1 - a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right). \tag{15}$$

Note that a factor of $\sqrt{2}$ appears in Equation 15 due to l_2-distance between x and x_i (Cauchy-Schwarz inequality) [27]. Combining Equations 10, 13, and 15 using triangle inequality results in,

$$Pr\left(\left|\hat{f}_{x|X} - f_x\right| \leq \left(\mathcal{L} + \sqrt{2}\mathcal{L}_1\right) d(x, x_i) + \frac{k\omega}{\sqrt{n + \frac{\omega^2}{\sigma_0^2}}}\right) \geq \text{erf}\left(\frac{k}{\sqrt{2}}\right)\left(1 - a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right)\right), \tag{16}$$

for any positive constant k. Note that Equation 16 holds for any measurement location x_i, which is a distance $d(x, x_i)$ away from the test location x and can be satisfied by choosing $d(x, x_i)$ stated in Lemma 2 for a given δ. □

Lemma 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 ϵ with high probability δ in a smaller disk of radius $\frac{1}{\alpha}r_{max}$ around x_i, where $\alpha > 1$ (Figure 1).

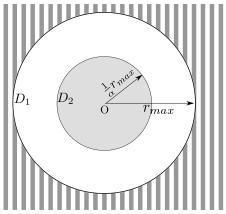

Fig. 1. Collecting n_α measurements at O suffices to make accurate predictions at all points inside disk D_2 (Sufficient condition). No number of measurements at O can ensure predictive accuracy on points outside disk D_1 (Necessary condition).

Lemma 3. *Given a disk of radius $\frac{1}{\alpha}r_{max}$ centered at x_i, with $\alpha > 1$, with high probability δ, n_α measurements at x_i suffice to make ϵ-accurate predictions for*

all points inside the disk, where,

$$n_\alpha \geq \left(\frac{\sqrt{2}\omega \operatorname{erf}^{-1}\left(\frac{\delta}{1-a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right)} \right)}{\epsilon - \frac{l\left(\mathcal{L}+\sqrt{2}\mathcal{L}_1\right)}{\alpha}\sqrt{-\log\left(\frac{\sigma_0^2+\omega^2}{N\sigma_0^2}\left(1-\frac{\epsilon}{\sqrt{2}\sigma_0^2\operatorname{erf}^{-1}(\delta)}\right)\right)}} \right)^2 - \frac{\omega^2}{\sigma_0^2}. \tag{17}$$

Proof. We want a sufficiency condition on the number of measurements n_α inside a disk of radius $\frac{1}{\alpha}r_{max}$. Lemma 2 gives an upper bound on the radius of a disk such that all points inside the disk will be chance-constrained after n_α measurements at the center. We construct a disk (D_2 in Figure 1), whose radius is equal to $\frac{1}{\alpha}r_{max}$ such that,

$$\frac{1}{\alpha}r_{max} \leq \frac{1}{\mathcal{L}+\sqrt{2}\mathcal{L}_1}\left(\epsilon - \sqrt{\frac{2}{\frac{1}{\sigma_0^2}+\frac{n_\alpha}{\omega^2}}}\operatorname{erf}^{-1}\left(\frac{\delta}{1-a\exp\left(-\frac{\mathcal{L}_1^2}{b^2}\right)}\right)\right). \tag{18}$$

This ensures all points in D_2 will be chance-constrained. Substituting the value of r_{max} from Lemma 1 in Equation 18 and and re-arranging for n_α gives the required bound stated in Lemma 3. □

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 $\frac{1}{\alpha}r_{max}$.

4.3 Placement of Sensors for Problem 1

We use an algorithm similar to the one presented by Tekdas and Isler [28]. The exact procedure is outlined in Algorithm 1. We choose $\alpha = 2$.

Algorithm 1: DiskCover

input : An arbitrary environment
output: Measurement locations

1 Given an arbitrary set $\mathcal{X}$ of disks covering the given environment, calculate the Maximal Independent Set (MIS) $\mathcal{I}$ of $\mathcal{X}$ greedily *i.e.,* $\mathcal{I} = \text{MIS}(\mathcal{X})$;
2 Place disks of radii $2r_{max}$ at the center of each disk in $\mathcal{I}$. Call the set of $2r_{max}$ radii disks $\bar{\mathcal{X}}$;
3 Cover each disk in $\bar{\mathcal{X}}$ with disks of radii $\frac{1}{2}r_{max}$ as shown in Figure 2 and label centers of all disks of radii $\frac{1}{2}r_{max}$;
4 Return all the labeled points in previous step as measurement locations;

Theorem 1. DiskCover *(Algorithm 1) gives a* 32*-approximation for Problem 1 in polynomial time.*

Proof. Denote the optimal set of measurement locations to solve 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 any disk in $\mathcal{X}\backslash\mathcal{I}$ intersects with some disks in $\mathcal{X}$ (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 have to collect measurements from as many measurement locations as the cardinality of $\mathcal{I}$. This can be proved by contradiction. Suppose an algorithm visits less measurement locations than 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 1, the robot can never make accurate predictions at that point and hence violating the constraint in Problem 1. Hence,

$$|\mathcal{I}| \;\leq\; |X^*|. \tag{19}$$

The union of $\bar{\mathcal{X}}$ disks covers the entire environment. Consider a test point x in the environment and assume that it does not lie in the union of $\bar{\mathcal{X}}$ disks. Since no disk in $\bar{\mathcal{X}}$ contains x, x lies at least $2r_{max}$ distance away from the center of each disk in $\bar{\mathcal{X}}$ (also the centers of disks in $\mathcal{I}$). Hence, the nearest disk to x in $\mathcal{I}$ will be at least r_{max} away from x, as a result of which a disk of radius r_{max} around x can be constructed without intersecting any disk in $\mathcal{I}$ which violates the fact that $\mathcal{I}$ is an MIS.

Collecting measurements from 32 locations inside a $2r_{max}$ disk suffice to make accurate predictions in that disk (satisfying the Problem 1 constraint for points lying in that disk) as illustrated in Figure 2. DiskCover collects measurement from 32 such locations per disk in $\bar{\mathcal{X}}$. It collects measurements from a total of $32|\bar{\mathcal{X}}|$ 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. Using Equation 19, $32|\mathcal{I}| \leq 32|X^*|$. Note that $|\bar{\mathcal{X}}| = |\mathcal{I}|$. Hence,

$$32|\bar{\mathcal{X}}| \;\leq\; 32|X^*|, \tag{20}$$

$$n_{\textsc{DiskCover}} \;\leq\; 32|X^*|, \tag{21}$$

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

4.4 Finding 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 visited by the robot. The robot must obtain at least n_2 samples (Lemma 3 with $\alpha = 2$) at the center of each disk of radius $\frac{1}{2}r_{max}$. A pseudo-code of the algorithm is presented in the Algorithm 2.

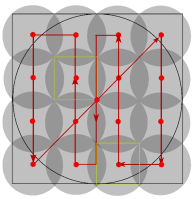

Fig. 2. We cover the disk of radius $2r_{max}$ using disks of radii $\frac{1}{2}r_{max}$ in lawn mover pattern which requires no more than double the minimum number of $\frac{1}{2}r_{max}$ disks required to cover a $2r_{max}$ disk. 32 disks suffice to cover the bigger disk whereas the minimum number of $\frac{1}{2}r_{max}$ disks required is $\frac{\pi(2\ r_{max})^2}{\pi(0.5\ r_{max})^2} = 16$. The locations of disks of radii $\frac{1}{2}r_{max}$ inside disk of radius $2r_{max}$ are obtained by covering a square circumscribing bigger disk with smaller squares inscribed in smaller disks. The centers of smaller squares coincide with centers of smaller disks.

Algorithm 2: DiskCoverTour

input : A set of measurement locations calculated from Algorithm 1.
output: Optimal path visiting all measurement locations.

1 Cover each $\bar{\mathcal{X}}$ disk in lawn-mover pattern visiting the centers of corresponding disks of radius $\frac{1}{2}r_{max}$ and make n_2 measurements at each center point;
2 Calculate approximate-polynomial TSP tour visiting the centers of $\bar{\mathcal{X}}$ disks;
3 Return union of paths calculated in steps above;

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

Proof. From Theorem 1 we already have a constant approximation bound on number of measurement locations. Let the time (travel and measurement time) taken by optimal algorithm be T^*. Using notation from Theorem 1, we assume that optimal TSPN time to visit disks in $\mathcal{I}$ be $T^*_{\mathcal{I}}$. 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 ($T^*_{measure}$),

$$T^*_{\mathcal{I}} \ \leq \ T^*_{travel}, \ \eta|\mathcal{I}| \ \leq \ T^*_{measure}. \tag{22}$$

Let denote optimal time to visit the centers of disks in $\mathcal{I}$ by $T^*_{\mathcal{I}^C}$. The time $T^*_{\mathcal{I}^C}$ to visit centers can be bounded by a path length having maximum detour of $|\mathcal{I}| \times 2r_{max}$ from $T^*_{\mathcal{I}}$. As a result: $T^*_{\mathcal{I}^C} \leq T^*_{\mathcal{I}} + 2r_{max}|\mathcal{I}|$. Using inequality from Equation 22: $T^*_{\mathcal{I}^C} \leq T^*_{travel} + 2r_{max}|\mathcal{I}|$. For any disk in $\bar{\mathcal{X}}$, the length of lawn-mover path starting from the its center and return back (Figure 2) after visiting all center points of $\frac{1}{2}r_{max}$ disks is $40r_{max}$. Hence, the total travel time for DiskCoverTour is: $T_C + 40r_{max}|\mathcal{I}|$, 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 [29] 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. Total measurement time for DiskCoverTour is $32\eta n_2|\mathcal{I}|$. Hence, total time T_{alg} for

DiskCoverTour is,

$$T_{alg} = T_C + 40r_{max}|\mathcal{I}| + 32\eta n_2|\mathcal{I}| \leq (1+\epsilon)\ T^*_{\mathcal{I}^C} + 40r_{max}|\mathcal{I}| + 32\eta n_2|\mathcal{I}|, \quad (23)$$

$$\leq (1+\epsilon)\ \ (T^*_{travel} + 2r_{max}|\mathcal{I}|) + 40r_{max}|\mathcal{I}| + 32\eta n_2|\mathcal{I}|. \quad (24)$$

Length of any tour that visits k non-overlapping equal size disks of radius r is at least $0.24kr$ [30], which gives $0.24r_{max}|\mathcal{I}| \leq T^*_{\mathcal{I}}$. Combining this result with Equation 22 modifies the bounds in Equation 24 as,

$$T_{alg} \leq \left((1+\epsilon)\left(1+\frac{2}{.24}\right) + \frac{40}{.24}\right) T^*_{travel} + 32n_2 T^*_{measure}, \quad (25)$$

$$\leq \max\left(9.33(1+\epsilon) + \frac{40}{.24}, 32n_2\right)\left(T^*_{travel} + T^*_{measure}\right), \quad (26)$$

$$\leq \max\left(9.33(1+\epsilon) + \frac{40}{.24}, 32n_2\right) T^* \leq cT^*, \quad (27)$$

where c denotes the corresponding maximum value in the right hand side of Equation 26, and is a constant. □

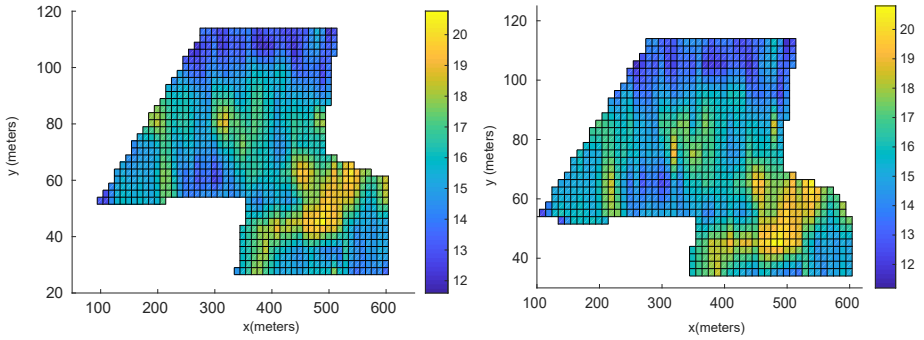

(a) The actual OM content. (b) Posterior prediction of OM content.

Fig. 3. Actual and predicted OM content comparison.

5 Simulations

We evaluate our algorithm with other baseline strategies through simulations based on a real-world dataset [31] of a farm (Figure 3(a)). The dataset [31] consists of organic matter (OM) measurements collected from 350 locations manually in a farm. Since the algorithm may require to obtain measurements at locations other than one of original 350 locations, we use cubic interpolation to generate the ground truth values at any location in the environment. The simulated sensor returns a noisy version of this ground truth measurement with an additive Gaussian noise of variance $\omega^2 = 0.25$.

We compare the performance of our algorithm against mutual information (MI) and entropy-based planning strategies [1] using three criteria. The squared-exponential kernel has two hyperparameters: namely length scale (l) and signal variance (σ_0^2). The values of l and σ_0^2 were estimated by minimizing the negative log-marginal likelihood for 350 locations where we know the OM content. We discretized the farm with a resolution of $\frac{l}{3}$ and calculated the values of required metrics at discretized locations. We used $\delta = 0.8,\ \epsilon = 2,\ \omega^2 = 0.25$ and l was estimated to be equal to 12.24 *meters* by log-marginal likelihood minimization. Figure 3(b) presents the predicted OM content in the farm by DiskCover algorithm. Figure 4 shows the covering of given environment using disks of radii $2r_{max}$ concentric with disks in the MIS calculated by DiskCover algorithm.

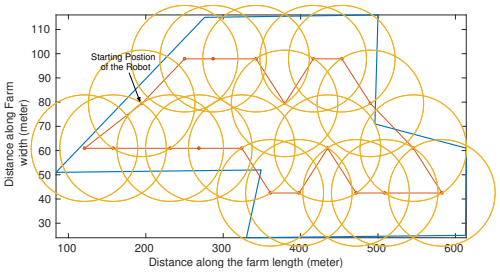

Fig. 4. Disk placement covering the farm calculated by *DC* algorithm. The disks are of radii $2r_{max}$ which are concentric with disks of radii r_{max} in $\mathcal{I}$. The optimal TSP tour visiting the centers and farm boundaries are shown in red and blue respectively. Lawn mover detours have been omitted to make the figure more legible.

The total time taken by the robot is summation of the travel time and the time taken to perform measurements. For MI and entropy-based planning, we calculate the measurement locations greedily [1]. After calculating the measurement locations for both, we find the optimal TSP tour visiting all the measurement locations and calculate the total time subsequently.

Figure 5 presents the percentage of points with prediction error more than 2 *ppm* as a function of time spent by the robot in the farm. DiskCoverTour outperforms MI and entropy-based planning in the beginning but as robot spends more time in the farm, MI and entropy-based planning perform better than DiskCoverTour. Asymptotically DiskCoverTour converges with fewer points (than MI and entropy-based planning) which have high prediction error. DiskCoverTour getting outperformed by other strategies in the middle can be explained by the observation that en-

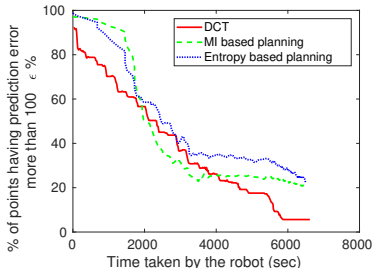

Fig. 5. Percentage of points with prediction error more than the maximum allowed error (2 *ppm*) as a function of tour time.

tropy and MI based planners tend to spread the measurement locations as far as possible from each other, thereby, covering a large part of the environment quicker than DiskCoverTour. A plot of Root Mean Square Error (RMSE) vs time is shown in Figure 6(a). We observe a similar trend as Figure 5. This trend can be attributed to the fact that MI and entropy-based planning methods are indirect methods and they tend to collect the measurements from locations which are far away from each other.

We plot entropy (a measure of posterior variance) to quantify the confidence about learned OM content variation. For our choice of $\delta = 0.8$, the entropy for DiskCoverTour algorithm converges to 230 as shown in Figure 6(b). For $\delta \approx 1$, it converges to the same values as MI and entropy-based strategies. Entropy-based planning achieves the maximum reduction in entropy of the distribution because of its diminishing returns property [1]. However, a reduction in an indirect measure like entropy does not guarantee reduction in prediction error. For this fact, entropy and MI-based strategies may cease to perform better than DiskCoverTour while achieving accuracy in posterior predictions.

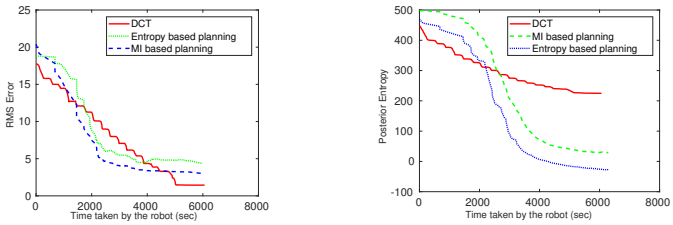

(a) Comparison of total error between actual organic matter content and predicted content.

(b) Entropy estimate for posterior distribution.

Fig. 6. Predictive error and posterior entropy comparisons.

6 Conclusion

We study the problem of minimizing the time taken by a single traveling robot to learn a spatial field. We provide a lower bound on the number of measurement locations required to estimate the field sufficiently accurate with high confidence. We also provide a polynomial time approximation algorithm whose performance is within a constant factor of the lower bound. 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. Our ongoing work is on developing competitive strategies for such cases as well as for spatio-temporal learning. As an immediate future work, we aim to remove the dependence of n_α on N and optimize for the value of α to improve the approximation ratio.

References

1. 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.
2. 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.
3. 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.
4. Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.
5. 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.
6. 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.
7. C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
8. 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.
9. L. Blackmore, M. Ono, and B. C. Williams. Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics*, 27(6):1080–1094, Dec 2011.
10. 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.
11. 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.
12. William F Caselton and James V Zidek. Optimal monitoring network designs. *Statistics & Probability Letters*, 2(4):223–227, 1984.
13. Dale L Zimmerman. Optimal network design for spatial prediction, covariance parameter estimation, and empirical prediction. *Environmetrics*, 17(6):635–652, 2006.
14. 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.
15. George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functionsi. *Mathematical Programming*, 14(1):265–294, 1978.

16. 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.
17. Dezhen Song, Chang-Young Kim, and Jingang Yi. Simultaneous localization of multiple unknown and transient radio sources using a mobile robot. *IEEE Transactions on Robotics*, 28(3):668–680, 2012.
18. Michael Otte, Nikolaus Correll, and Emilio Frazzoli. Navigation with foraging. In *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*, pages 3150–3157. IEEE, 2013.
19. Alexandra Meliou, Andreas Krause, Carlos Guestrin, and Joseph M. Hellerstein. Nonmyopic 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.
20. Y. T. Tan, Abhinav Kunapareddy, and Marin Kobilarov. Gaussian process adaptive sampling using the cross-entropy method for environmental sensing and monitoring. *IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia*, 2018.
21. 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.
22. 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.
23. 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.
24. 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.
25. Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2012.
26. Mike Giles. Approximating the erfinv function. In *GPU Computing Gems Jade Edition*, pages 109–116. Elsevier, 2011.
27. Sheldon Jay Axler. *Linear algebra done right*, volume 2. Springer, 1997.
28. Onur Tekdas and Volkan Isler. Sensor placement for triangulation-based localization. *IEEE transactions on Automation Science and Engineering*, 7(3):681–685, 2010.
29. 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.
30. 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.
31. D. J. Mulla, A. C. Sekely, and M. Beatty. Evaluation of remote sensing and targeted soil sampling for variable rate application of nitrogen. pages 1–15, Madison, USA, 2000. American Society of Agronomy.