

Title	Multi-UAV Resource Constrained Online Monitoring of Large-scale Spatio-temporal Environment with Homing Guarantee
Author(s)	Tiwari, Kshitij; Jeong, Sungmoon; Chong, Nak Young
Citation	43rd Annual Conference of the IEEE Industrial Electronics Society, IECON 2017: 5893-5900
Issue Date	2017-10-29
Type	Conference Paper
Text version	author
URL	http://hdl.handle.net/10119/15259
Rights	This is the author's version of the work. Copyright (C) 2017 IEEE. 43rd Annual Conference of the IEEE Industrial Electronics Society, IECON 2017, 2017, 5893-5900. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Description	

Multi-UAV Resource Constrained Online Monitoring of Large-scale Spatio-temporal Environment with Homing Guarantee

Kshitij Tiwari*, Sungmoon Jeong and Nak Young Chong

School of Information Science, JAIST, Japan

{kshitij.tiwari*, jeongsm, nakyoung}@jaist.ac.jp

Abstract—We propose a homing constrained bi-objective optimization variant of budget-limited informative path planning for monitoring a spatio-temporal environment. The objective function consists of weighted combination of two components: *model performance* which must be maximized and *travel distance* which must be bounded by the maximum operational range. Besides this, we have additional constraints that guarantee that the robots will return to home (base station) upon completion of their respective missions. Optimizing over this objective function is essentially NP-hard owing to the conflicting constituents. Moreover, the appropriate choice of weights and additional homing guarantees further adds to complications. We employ Gaussian Process (GP) model [1] which is highly data driven i.e., the larger the amount of training data, the better the model performance. However, owing to limited resources, a robot can only collect a limited amount of training samples. Thus, with the introduction of our bi-objective cost function, it becomes possible to plan budget-limited (e.g., battery, flight time, travel distance etc.) informative tours using autonomous mobile robots to effectively select only the most informative (uncertain) locations from the environment. In this work, we develop an algorithm to autonomously choose the appropriate weights for the components based on available resources while ensuring homing and maintaining model quality. We perform simulations to verify the effectiveness of our proposed objective function on the publicly available Ozone Concentration dataset gathered from USA.

I. INTRODUCTION

Monitoring and modeling a large scale spatio-temporal phenomenon using mobile robots with limited resources (e.g., battery life, travel distance, payload capacity) necessitates efficient utilization of resources whilst not compromising on the model performance. Monitoring spatial variations and temporal dynamics of any indoor/outdoor spatio-temporal phenomenon like indoor room temperature [2], environmental ozone monitoring [3], [4], oceanic phytoplankton density monitoring [5], daily precipitation monitoring [6], traffic flow density for modeling mobility-on-demand (MoD) patterns [7] require large amounts of data to be gathered and processed. Usually for processing the data, a non-parametric Bayesian framework called Gaussian Process (GP) [1] is a preferred model of choice. Besides elegantly capturing the complex underlying dynamics of the spatio-temporal environment, these models also give a theoretical measure of model uncertainty which has proven to be useful to evaluate the model quality. Since the robot has limited resources, gathering informative input data for training the model is critical. Thus, optimal input selection under resource constraints or *efficient active*

sensing for intelligent environment monitoring will be the main focus of this paper. We aim to address the following research question: *Given two inherently conflicting objectives viz., model performance and resource utilization, how can we effectively optimize over both of them (simultaneously) while ensuring homing?*

GPs are highly data driven models and hence, the choice of training samples largely affects the model performance. Usually, two kinds of criteria are used for active sensing: *Firstly, entropy maximization* whereby the absolute entropy over the unvisited area is considered to deduce the most uncertain location and *Secondly, mutual information gain* criteria which evaluates the reduction in entropy obtained if a candidate location was actually observed. The limitation of *entropy maximization* is that it forces the robot to move to locations which are prohibitively far away incurring huge costs (sensing cost, travel cost, battery) as explained in [3], [8]. An alternative to this is to use *mutual information* based active sensing. However, this is NP-complete [8] and hence, researchers tend to prefer polynomial time approximations instead. Another possibility could be to utilize infinite dimensional Kalman filters [9], however this would involve approximating the infinite dimensional stochastic differential equations and there is no prior works underlining how to account for resource constraints and homing.

Resource constrained robots have previously been considered for mapping and energy efficient path planning of unknown environments like in [10]–[13]. Information-theoretic path planning has been studied in [2], [6], [12], [14], [15] wherein the maximal gradient of information gain is followed by robots to actively gather the training samples. The benefit of using information-theoretic path planning is that it can be easily coupled with the model being used to generate the estimations over the target environment to drive the robots to autonomously gather the best of the available locations. In [5], [7] the authors use the maximum sampling budget as a condition for termination of exploration. However, none of these approaches guarantee “homing”, i.e., ensuring that the robots return to the base station upon exhaustion of available budget or reaching a termination condition. Also, while ensuring “homing” we need to ensure that the model quality is not compromised. Similar problems with homing have been considered in [16], [17] where the researchers pose this problem as a variant of *Orienteering problem*. However, in these works, we cannot concurrently learn the

model parameters and estimate the measurements.

A very naïve solution to ensure thrifty resource utilization could be to minimize the resources lost while returning home. However in doing so, the robot will try to return home as early as possible which means it will gather minimal training samples and would greatly compromise the model performance. To overcome such a situation and elegantly trade-off model performance to resource utilization, in our prior work [3], we proposed a bi-objective optimization problem. In doing so, we combined two conflicting objectives into a single cost function using a weighted linear combination but the weights of the constituent objectives were arbitrarily chosen. After further analysis, we noted that the choice of weights can affect the quality of training samples which in turn affects the model performance. Only focusing on *model performance* would force the robot to gather high quality measurements which are few in number, whereas putting a high weight on *travel distance* would force the robot to visit only nearest neighbors which in a spatiotemporal setting have highly correlated measurements and hence do not add information to the model. Thus, carefully and correctly identifying the weights of the objective functions can be challenging [18].

We propose to (a) model the spatio-temporal environment using distributed GP framework from [3] and (b) extend the previously proposed Resource Constrained Decentralized Active Sensing (RC-DAS) from [3] wherein the robot trades-off the model performance to travel distance. This extension now ensures that the robot always returns back to base station (home) whilst not compromising on the performance significantly. Homing is essential not only for prevention from losing robots in the middle of the field but also for the next phase of our architecture where we will use the learnt models to fuse them into a globally consistent model at the base station.

II. PROBLEM STATEMENT

In this section, we formally introduce the sensing scenario and the overall system architecture.

A. Sensing Scenario

We consider a d dimensional sensing domain $D \subset \mathbb{R}^d$ represented as a network of spatio-temporally correlated nodes or pre-determined locations like [16], [17]. Such a scenario is to be modeled using a mobile robot which can only obtain measurements at these a priori known locations under resource limitations (like battery life, travel distance, etc.). The robot behaves like a self-sustaining GP expert, always moving deterministically to gather its training samples and self localizing with respect to the map being generated. However, in doing so, the robot must always inspect the available resources and ensure that at the end of the tour ¹, it can return home as is shown in Fig. 1.

¹We define a tour as a path traced (sequence of states attained by the robot) from an arbitrarily assigned start location (selected from one of the nodes) and terminating at the base station (the location of which is also known a priori to the robot).

From the Figure, we can see that there are two alternatives: *Firstly*, the robot can visit the most uncertain (informative) location with no homing guarantees or *Secondly*, compromise on the quality of the model by observing a closeby region instead, with homing guarantees. The latter alternative is preferable and will be the main focus of this research. Given a team of multiple UAV's with homing constraints, upon termination of missions of all robots, we can easily fuse the multiple decentralized models generated by the robots into a globally consistent model as shown in Fig. 3. The advantage of posing our architecture like this is that, each robot can independently control its own path planning and enhance its model quality to the best of its capability while the base station can generate the final model for which the only requirement is for all robots to return home so that the base has access to their respective models.

Example 1 (Sample Application): The reason for such a choice of constraints becomes comprehensible from the following example: Consider a nuclear disaster situation. Since the environment is toxic, usually the humans workers and base stations are positioned at a safe distance from the disaster site. Then the robots are set free to explore and infer the situation. Now consider this as a multi robot setup since the area to be monitored is significantly large. In such a setting, we pose this as a decentralized sensing with centralized fusion algorithm. The reason being that each robot can choose to observe only a subset of the environment and generate a model. Later when all robots return to base station at the end of their mission times, we can fuse the individual models into a globally consistent model to generate a more precise model of the environment. Like this, we can gather sufficient amount of data whilst satisfying the resource constraints imposed on each member of the team.

III. PRELIMINARIES

In this section, we formally introduce the non-parametric Bayesian model called Gaussian Process (GP) and explain how inference is performed by carefully selecting the most important locations to be observed. The notational convention followed in this manuscript are as follows: All sets are represented by upper-case alphabets like D , all vectors are represented using bold lower-case characters/symbols like \mathbf{x} , all scalars are represented by regular lower-case characters like l_1 and all distributions are represented using the scripted fonts like \mathcal{GP} . $|\cdot|$ represents cardinality and $\|\cdot\|$ represents the Euclidean distance.

A. Gaussian Process (GP)

GPs are a rich class of non-parametric Bayesian models, which allow us to model spatio-temporal environmental phenomena. GPs consistently quantify the uncertainty associated with predictions (e.g., based on mean-squared error, entropy or mutual information gain criterion [5]), that can be exploited by active sensing schemes for exploration and obtaining the *most informative sensing locations* for each mobile robot. GPs belong to the family of Gaussian distributions and hence they can be easily characterized (by covariance

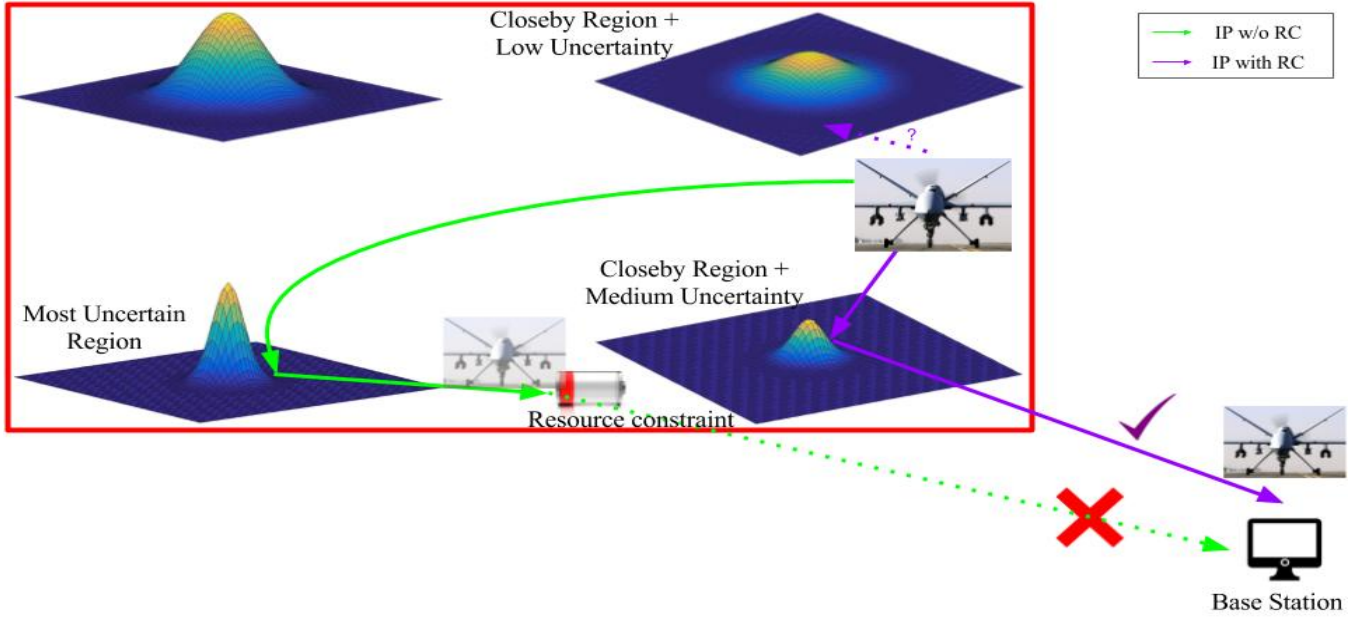


Fig. 1. (**Homing Scenario**) The robot utilizes Information-theoretic Path Planning (IP) for choosing the next best location to observe. Without any resource constraints (RC), it is biased towards visiting the most uncertain regions to reduce the uncertainty about the environment dynamics. Usually these locations/regions tend to be quite far away and sometimes, the robot may not have enough available resources to make it back to the base station. Thus, instead of losing our robots we prefer observing locations which may be slightly less informative but guarantee a return path to base station. The red bounding box represents the sensing limits and all locations that can be observed are pre-defined and a priori knowledge for the robot.

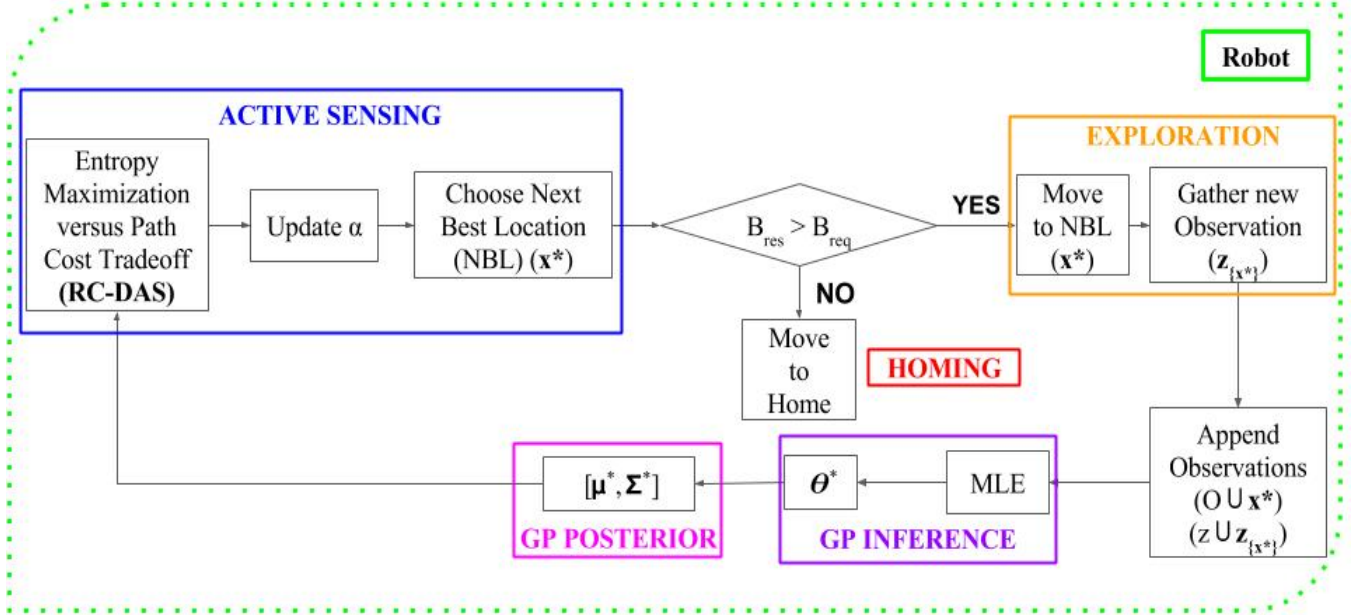


Fig. 2. (**System Architecture**) Here, we demonstrate the overall flow of our RC-DAS framework which serves the dual objective to trade-off the objective functions and terminate the exploration. In the **active sensing** block, we use our RC-DAS cost function to select the most informative location \mathbf{x}^* . If we have sufficient budget, then we move to \mathbf{x}^* for **exploration** and gathering the observation. The observed measurement $\mathbf{z}_{\mathbf{x}^*}$ is then stored and used to update the parameters of the GP model by re-performing MLE in the **GP inference** block. Upon completion, we now have access to updated **GP Posterior** which is used again for **active sensing** until a termination condition is reached which enforces **Homing**.

functions) and provide the uncertainty bounds associated with prediction estimates to evaluate model performance.

A GP is a generalization of a Gaussian distribution and fully defined by a mean function $\mu(\cdot) = \mathbb{E}[f(\cdot)]$ and covariance function $k(\cdot, \cdot)$. The covariance function (also

known as kernel), defines the spatio-temporal correlation structure of the function to be modeled and is parametrized by a set of hyper-parameters denoted by θ .

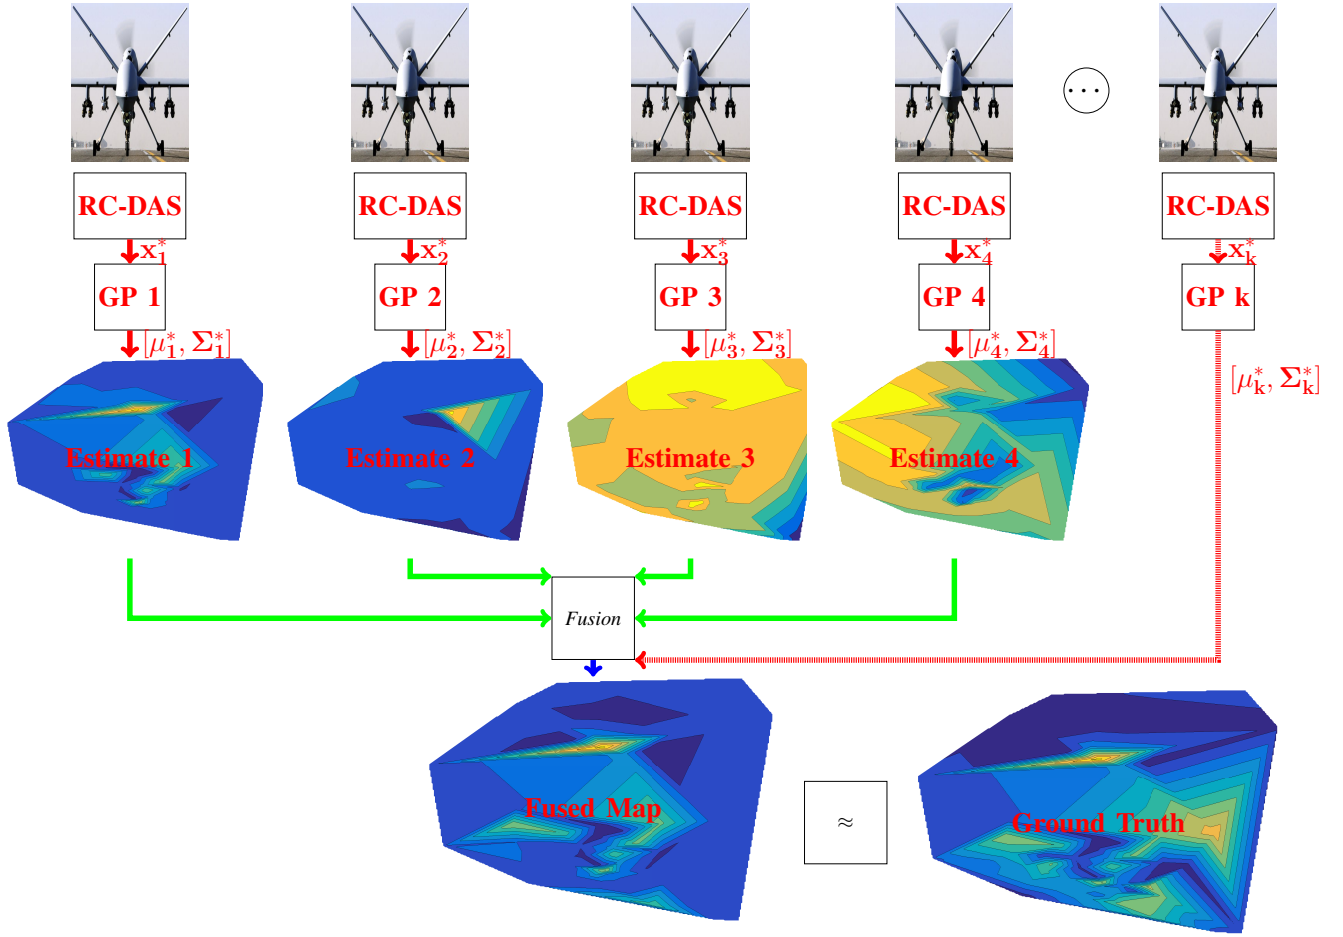


Fig. 3. **(Overall Sensing Scenario)** Illustration of the sensing scenario in which the team of mobile robots operates under resource constraints. The aim is to gather optimal observations to make a prediction for the environment defined by posterior mean μ_k^* and posterior covariance Σ_k^* . *Estimate 1–Estimate 4* represent the 4 individualistic prediction maps made by the 4 robots based on their training samples. \mathbf{x}_k^* represents the *next-best-location* chosen by the *RC-DAS* active sensing for k^{th} expert under homing constraints. Fused Map is the globally consistent fused prediction map generated by combining all individual models. Our target is to make the Fused Map as similar to the Ground Truth as possible. These maps have been interpolated for ease of visualization. In reality, we just have a discrete collection of predicted measurements at pre-determined locations since we have point sensing.

B. Inference in GP

A commonly used covariance function is the squared exponential covariance defined as:

$$k(\mathbf{x}, \mathbf{x}') = \sigma_{sig}^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T L^{-1}(\mathbf{x} - \mathbf{x}')\right) + \Sigma_n \quad (1)$$

where $\mathbf{x}, \mathbf{x}' \in D$, $L = \text{diag}(l_1, \dots, l_d)$ and the l_i are characteristic length scales at state t , which determine the relevance of the corresponding input dimension for modeling the spatio-temporal phenomenon. σ_{sig} corresponds to the amplitude of the signal to be modeled whereas Σ_n describes the magnitude of the noise. The hyper-parameters are $\theta \triangleq \{\sigma_{sig}, \Sigma_n, l_1, l_2, \dots, l_d\}$. The hyper-parameters are trained using the standard procedure of evidence (type-II marginal likelihood) maximization [1]. Evidence maximization avoids over fitting by automatically trading off data fit and model complexity. In our multi-robot setting, for each robot, we define $O \subset D$ as the set of observed nodes in the spatio-temporally correlated domain D and $U \subset D$ as the set

of unobserved inputs such that $U = O^c$. We also define $\mathbf{z}_{\mathbf{x}}; \forall \mathbf{x} \in O$ as the observed measurements and $\mathbf{z}_{\mathbf{x}^*}; \forall \mathbf{x}^* \in U$ as the corresponding predictions at the unobserved locations. The posterior mean and covariance of a GP are given by:

$$\mu_{U|O, \theta} \triangleq \mu_U + \Sigma_{UO|O} \Sigma_{OO|O}^{-1} (\mathbf{z}_O - \mu_O) \quad (2)$$

$$\Sigma_{UU|O, \theta} \triangleq \Sigma_{UU|O} - \Sigma_{UO|O} \Sigma_{OO|O}^{-1} \Sigma_{OU|O} \quad (3)$$

C. Active sensing using GPs

As explained in [3], we will use the posterior entropy as a measure of uncertainty for the GP models to deduce the most informative *next-best-location*. Thus, we define:

$$\mathbb{H}_{\mathbf{z}_U | \mathbf{z}_O, \theta} \triangleq \frac{|U|}{2} \ln(2\pi e) + \frac{1}{2} \ln(|\Sigma_{UU|O, \theta}|) \quad (4)$$

In Eq. (4), $\mathbb{H}_{\mathbf{z}_U | \mathbf{z}_O, \theta}$ refers to the entropy over the posterior covariance $\Sigma_{UU|O, \theta}$ of size $|U| \times |U|$.

IV. PROPOSED FRAMEWORK

In this section, we explain our proposed system architecture that ensures homing whilst optimizing model performance. To achieve this, we extend our previous work to propose a constrained bi-optimization variant which is also detailed herein.

A. System Architecture

Our system architecture is shown in Fig. 2: The first step while performing active sensing using our *RC-DAS* approach is to decide the weight (α) for the objective functions, that will trade off resource utilization to model performance. In essence, our cost function looks like $(\alpha \text{ Entropy} + (1 - \alpha) \text{ Distance})$ wherein we want to emphasize the most on the first term when the resources are sufficient to ensure the model performance is enhanced and gradually shift focus to the second term as the resources decay to ensure safe return to base. Once we have deduced this weight metric, we also check, if the current available budget B_{res} is larger than the required budget B_{req}^2 to be able to reach the home base via the next location. If not, we terminate the robot exploration. However, if we have sufficient residual budget to be able to cover the net cost to return to base station from \mathbf{x}^* , we choose to move to \mathbf{x}^* . Upon reaching the next best location, \mathbf{x}^* , we collect the observation $\mathbf{z}_{\mathbf{x}^*}$ and store it along with the input location \mathbf{x}^* to update the model parameters, thereby obtaining the updated posterior mean $\boldsymbol{\mu}^*$ and covariance Σ^* . We then use the posterior covariance for active sensing.

B. Framework

Other works like [19] also consider a principled approach for multi-agent team for online inference of GP model but they assume that communication channels are perfect and always connected. In contrast to this idealistic setting, we propose our model for a rather harsh communication-devoid environment. This paper serves to further enhance and address the limitations of [16]³:

- **Choosing the next-best-location:** Instead of looking at just the immediate neighbors, our approach evaluates all correlated locations within the field w.r.t. the current location (\mathbf{x}) to deduce the most informative candidate.
- **Concurrent Online Inference and Estimation:** In our approach, we can concurrently *infer* the optimal model parameters by updating the model as and when new data comes in and if required, *estimate* the measurement at any arbitrary input location.
- **Measurement Noise:** All measurements gathered are considered to be noisy and the noise variance is itself treated as a parameter to be learnt via inference as opposed to Yu et.al's work where noise free measurements were considered.
- **Informativeness of a location:** In [16], the informativeness of a candidate location j was only considered

with respect to a specific location i in its immediate neighborhood, independent of the rest. However, in our case, we evaluate the informativeness of a candidate location in terms of the reduction of uncertainty achieved over the entire environment (i.e., all the unobserved locations).

- **Homing Guarantees:** Similar to [16], our cost function explicitly guarantees that at the end of the mission time each robot can safely return to base station and will not get stranded amidst the field. However, in [16], this was done offline while we dynamically adjust the exploration strategy based on available resources.
- **Scaling to multiple robots:** Our architecture is easily scalable to a fully decentralized robot team where each agent is individually optimizing its own resources and model performance to generate the best feasible model under homing guarantees. In doing so, we ensure that no robot gets stranded and we can generate a globally consistent model of the environment based on several individual models generated by the team as illustrated in Fig. 3.

These extensions are proposed under the following assumptions:

- The resources (battery life, flight time, travel distance, etc.) available to the robot are not enough to perform exhaustive coverage of the phenomenon. Blanket coverage may lead to supreme model performance but is practically infeasible⁴ and the robot is tasked with planning budget-limited informative tour to myopically maximize the reward (measured in terms of information gain).
- The tour of every robot is assumed to have concluded within each *time step*⁵.
- Most of the environment monitoring datasets only record measurements for static sensors placed at discrete locations. However, not all stations need to be observed at all times. In this work, we aim to select the “key” locations to be observed. Thus, we no longer have access to static sensors but since the measurements are only available at the locations where the static sensors were previously placed, we restrict our robots to only observe and visit these locations. Thus, in our setting, the locations that can be observed are pre-defined and known to robots a priori.⁶ (similar to [16]).

Befitting these assumptions, we wish to model a spatio-temporal environment $\mathbf{z} = f(\mathbf{x}) + \epsilon$, where $\mathbf{x} \in D$ are inputs and $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma_n)$, where $\Sigma_n = \text{diag}(\sigma_{l_1}, \dots, \sigma_{l_d})$, is i.i.d. Gaussian measurement noise. We place a Gaussian process (GP) prior on the spatio-temporal phenomenon f and write

⁴the limitation not only arises owing to limited resources but also owing to point sensing as opposed to range sensing, i.e. our sensors have measurements only at the current location instead of a region.

⁵*time step* refers to the quantum in temporal domain.

⁶It must be noted here that GPs can be used to predict measurements at any arbitrary location but since the ground truth for such locations was not made available in the raw dataset, we cannot evaluate the prediction performance and hence were not considered in the current problem setup.

² $B_{req} \triangleq C_S(\mathbf{x}^*) + C_T(\mathbf{x}, \mathbf{x}^*) + C_T(\mathbf{x}^*, \text{Home})$

³Since, not all our model assumptions are satisfied by the said work, we do not include results from [16] for empirical analysis.

$f \sim \mathcal{GP}(\boldsymbol{\mu}(\cdot), k(\cdot, \cdot))$. We associate a sensing cost $C_S(\mathbf{x}^*)$ and a travel cost $C_T(\mathbf{x}, \mathbf{x}^*)$ for observing measurements at new location \mathbf{x}^* . We define a maximum budget as B and the remaining available budget as B_{res} .

C. Variants of Decentralized Active Sensing (DAS)

In this section, we briefly recapitulate the two variants of decentralized active sensing being considered in this work. These have already been explained in depth in [3]. Besides this, we also extend our previously proposed active sensing scheme to enforce homing.

1) *Fully Decentralized Active Sensing (full-DAS)*: For *full-DAS*, the main aim is to visit the most uncertain locations as long as the robot has not run out of resources (battery). Thus, the cost function is formulated as:

$$\hat{\mathbf{x}}_F = \arg \max_{\mathbf{x}^*} \mathbb{H}[\mathbf{z}_{U_{\mathbf{x}^*}} | \mathbf{z}_{O_{\mathbf{x}}}] . \quad (5)$$

where it can be noted from (5) that cost function cannot guarantee that the robot can safely return to base, since only the model performance is being optimized.

2) *Resource Constrained DAS (RC-DAS)[†]*: As opposed to *full-DAS*, in our *RC-DAS* framework [3], the cost function comprises of two objective functions: *entropy* to be maximized and *travel distance* to be minimized. These two objectives are highly conflicting and optimizing them simultaneously is rather challenging. Our new homing constrained bi-objective *RC-DAS* cost function is given by:

$$\begin{aligned} \hat{\mathbf{x}}_R = \arg \max_{\mathbf{x}^*} & (\alpha \mathbb{H}[\mathbf{z}_{U_{\mathbf{x}^*}} | \mathbf{z}_{O_{\mathbf{x}}}] - (1 - \alpha) \ln \|\mathbf{x} - \mathbf{x}^*\|) , \\ \text{s.t. } \arg \min_{\mathbf{x}^*} & \{C_T(\mathbf{x}, \text{Home}), C_T(\mathbf{x}^*, \text{Home})\} \end{aligned} \quad (6)$$

where $\mathbf{x} \in O$, $\mathbf{x}^* \in U$. In Eq. (6), if $\alpha = 1$, then the decentralized active sensing problem is similar to *entropy maximization* i.e. *full-DAS* as shown in Eq. (5). For $\alpha = 0$, the equation becomes similar to nearest neighbor selection routine without any active sensing. Similar cost functions have previously been used in [20], but the choice of α was made empirically. However, in this work, we propose a novel technique to choose the *weighting factor* (α) based on residual resources (B_{res}). The *state* (t) of the robot is characterized by $\{\mathbf{x}, B_{res}^{[t]}, \boldsymbol{\theta}^{[t]}\}$. For each new state, we define the weights as:

$$B_{res}^{[t]} \triangleq B_{res}^{[t-1]} - (C_S^{[t]}(\mathbf{x}^*) + C_T^{[t]}(\mathbf{x}, \mathbf{x}^*)) \quad (7)$$

$$\alpha^{[t]} \leftarrow \frac{B_{res}^{[t]}}{B} . \quad (8)$$

In Eq. (8), the current *weighting factor* ($\alpha^{[t]}$) is determined based on the current residual budget ($B_{res}^{[t]}$) as defined in Eq. (7). In Eq. (7), we define the current residual budget ($B_{res}^{[t]}$) as the difference between the previously available budget $B_{res}^{[t-1]}$ and the cost that was incurred in moving and gathering measurement in the previous time step ($t-1$) such that $0 < B_{res}^{[t]} \leq B$. At $t = 0$, we define $B_{res}^{[0]} \triangleq B$ such that $\alpha^{[0]} \leftarrow 1$. When plugging in apt weights (α) from Eq. (8)

into Eq. (6), we not only find the *next-best-location* but also check the condition if it is feasible to return home if the next location is actually visited, thereby guaranteeing homing.

V. EXPERIMENTS

In this section, we empirically evaluate the model performance using real world dataset. Since, our approach is fully decentralized and there is no communication between robots during exploration, we simply show the experiments for 1 robot case. Similar performance is expected irrespective of the size of the robot team since each robot independently optimizes its own model and is guaranteed to return to base at the end of its mission. Even though our team is fully decentralized we ascertain that they will not have completely overlapping tours. This is based on the fact that the robots starting at different locations will have different information gradients which is reflected their active sensing thereby reducing the overlap.

A. Datasets for performance evaluation

For empirical evaluation, we utilize the **US Ozone Dataset**. This dataset includes ozone concentrations (in parts per billion) collected by US Environmental Protection Agency [3]. In this dataset, the measurements were recorded for several years at 59 static monitoring stations across USA but we only choose one of the years for evaluation purposes. For each station, the annual average ozone concentration was assigned as the sample measurement for that station.

B. Experimental Setup

Since, inherently *full-DAS* does not ensure homing, we perform two sets of experiments: 1.) we enforce homing for *full-DAS* for a fixed budget. For doing so, we check if $C_S^{[t+1]}(\mathbf{x}^*) + C_T^{[t+1]}(\mathbf{x}, \mathbf{x}^*) + C_T^{[t+1]}(\mathbf{x}^*, \text{Home}) \leq B_{res}^{[t]}$ then we allow the robot to attain the next location \mathbf{x}^* otherwise we terminate the exploration. While this criterion is simply used to terminate exploration for *full-DAS* case, we utilize this condition to actually influence the choice of next best location for our proposed *RC-DAS[†]* framework. This provides a fair comparison to evaluate the relative improvement in the model quality when homing is enforced in both cases. *RC-DAS[†]* is expected to visit more locations for the same budget therefore improving the model performance. 2.) In the second set, we evaluate *full-DAS* without homing constraints to evaluate the probability that the robot can safely return to base station while *RC-DAS[†]* always guarantees homing; The motivation behind this is to evaluate the performance trade-off with homing of *RC-DAS[†]* to that of *full-DAS*. For each of these experiments we assign all of the 59 stations as possible start locations and analyze the results over all of them.

C. Evaluation Criterion

We use two kinds of evaluation criterion to illustrate the necessity for enforcing homing and the impact on model performance when homing is enforced. The model performance is defined as:

Definition 1 (Model Performance): The model performance when using a chosen active sensing scheme is defined as the *Root Mean Squared Error (RMSE)* over the predicted measurements for $\forall u \in U$ for a robot. Lower the RMSE, the better is the model performance and hence more accurate is the map.

Our evaluation criterion are defined below:

Definition 2 (Precision (P)): If a total of N experiments are performed during which N_F represents the number of times *full-DAS* generated a more accurate map than *RC-DAS* and N_R represents vice versa, then precision (P) for *full-DAS* is given by:

$$P_F \triangleq \frac{N_F}{N}, \quad (9)$$

and the precision (P) for *RC-DAS*[†] is given by:

$$P_R \triangleq \frac{N_R}{N}. \quad (10)$$

Thus, P represents the chances of generating a *better model*⁷ of the environment given the choice of active sensing scheme. We evaluate the accuracy of the model by comparing the predicted measurements with the ground truth values and evaluating the RMSE to associate a scalar value as a performance measure for the model being considered. This is strongly motivated by Experiment 1 to represent the expected trade-off of model performance to resource utilization.

Definition 3 (Uncertainty of Homing (UoH)): Given the start location of a robot x , the Uncertainty of Homing (UoH) represents the uncertainty in reaching the base station. If p_H is the probability of reaching the base station given x , then we define :

$$UoH \triangleq \neg p_H. \quad (11)$$

D. Results

In this section, we summarize our empirical analysis using the Ozone dataset.

1) *Uncertainty of Homing when no homing is enforced on full-DAS:* When no homing is enforced on *full-DAS*, the robot is purely focused on gathering as many informative training samples from the environment until the entire budget is exhausted. This could mean that the robot does not even have enough budget remaining to return to base station. In our experiments, while *RC-DAS*[†] guarantees $p_H = 100\%$ i.e. $UoH = 0\%$, for the *full-DAS*, the chances that the robot could return to base are $p_H \leq 1\%$ i.e. $UoH > 99\%$.

2) *Precision of RC-DAS[†] versus full-DAS in terms of prediction performance:* In this set of experiments, we evaluate the precision of *RC-DAS*[†] compared to *full-DAS*. We evaluate all possible start locations and report the average performance. Since *RC-DAS*[†] always considers homing, we performed two subsets of experiments: *Firstly*, considering *full-DAS* in its current form i.e. without homing and *Secondly*, by manually enforcing homing on *full-DAS*.

⁷c.f. Definition 1.

From Table I, we conclude that *full-DAS* has a higher precision when homing is not performed but owing to homing constraints our *RC-DAS*[†] has superior performance. Alternatively, this also tells us that when homing is a necessary condition, then our *RC-DAS*[†] is more robust to the choice of start location assigned to the robot. The start locations directly affect the trajectory and the terminal quality of the prediction model and hence robustness to the choice of start location is of utmost importance.

3) *Utilization of Budget and Length of Walk:* The number of observations gathered by a robot during its mission time is hereby referred to as the Length of Walk. From Table I, we already know that *RC-DAS* has superior performance when homing is enforced. To further support our claims, we also evaluate the budget utilization by both active sensing schemes as shown in Fig. 4 wherein each trend represents a single tour until budget exhaustion from a chosen start location.




Fig. 4. **(Budget Decay)** Analyzing how the budget is consumed (decayed) while gathering observations using the *full-DAS* and *RC-DAS*[†] active sensing schemes. Test are also performed for artificially enforced homing constraint for *full-DAS*. Each trend represents budget decay for the respective scheme for a chosen starting point.

From Fig. 4, we conclude that *RC-DAS*[†] is more conservative in utilizing the available resources and hence can allow the robots to observe more locations. The length of walk of *RC-DAS*[†] is significantly larger than *full-DAS*. When comparing *full-DAS* with and without homing, we can see that for most of the start locations, the length of walk for homing case is shorter than non-homing case.

TABLE I
IMPACT OF HOMING ON THE PRECISION OF *full-DAS* VS *RC-DAS*[†].

	P_F	P_R
Full-DAS w/o Homing	63.33%	36.67%
Full-DAS with Homing	45%	55%

4) *Experimental Synopsis*: In conclusion, we empirically ascertain that the precision of $RC-DAS^\dagger$ is significantly better than the state-of-the-art *full-DAS* variant. Our proposed framework allows conservative utilization of resources thereby accruing more observations while guaranteeing a safe return path to base station, thereby, satisfying its dual purpose.

VI. CONCLUSION AND FUTURE WORKS

The aim of this work was to pose decentralized active sensing for informative sampling as a bi-objective optimization problem where the objectives dominantly conflict each other. Under this setting, our cost function comprises of two objective functions: *measurement uncertainty* and *travel distance*. We called this approach as *Resource Constrained Decentralized Active Sensing (RC-DAS)[†]*, in which, as opposed to our works in [3], the robot now chooses the weights based on residual resources rather than pre-encoded heuristics. In doing so, we have the following benefits over our prior work [3]: 1.) the robot is guaranteed to return to base station at the end of its mission time. Thus, we not only ensure safe recovery of a robot at the end of its mission but this is also crucial for the next phase of our architecture where we pose this as a multi-robot problem such that we need to fuse prediction models from individual robots. This is done by the base station at the end of mission times of all robots and losing a robot could mean loss of mission-critical information. 2.) the cost function now is more robust to choice of start locations and 3.) the robots can dynamically trade-off resource utilization to model quality as a function of their available resources.

Our preliminary empirical analysis showed that when homing is enforced our $RC-DAS^\dagger$ outperforms the current state-of-the-art *full-DAS* approach but there is still scope of further improving our cost function. Ours is just a surrogate function which helps to apply multi objective optimization to informative path planning domain. In future, we will investigate the GP-LVM model [21] for efficient mobile robot localization to demonstrate our approach for real robots and so closing the exploration-loop could help enhance localization accuracy. Besides, we would also like to investigate partially decentralized approaches like [7] so that at the cost of minimal communication overhead we can co-ordinate our team and avoid overlapping observations. Integrating the robot dynamics model together with our cost function could also help to evaluate the feasibility of a location not only based on the available residual resources but also based on the robot dynamics. Further to this, uncertainty in resource utilization, robot motion, external disturbances (like wind etc.) and robot failures could be considered to incorporate robust planning over longer horizons.

ACKNOWLEDGMENTS

This work was supported by the Industrial Convergence Core Technology Development Program (No. 10063172) funded by MOTIE, Korea. The authors would like to thank

Mr. Anand Krishnamoorthy for constructive discussion that helped to enhance the quality of our work.

REFERENCES

- [1] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- [2] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin, "Efficient planning of informative paths for multiple robots," in *IJCAI*, 2007.
- [3] K. Tiwari, V. Honoré, S. Jeong, N. Y. Chong, and M. P. Deisenroth, "Resource-Constrained Decentralized Active Sensing for Multi-Robot Systems using Distributed Gaussian Processes," in *ICCA*, pp. 13–18, 2016.
- [4] R. Marchant and F. Ramos, "Bayesian optimisation for intelligent environmental monitoring," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*, pp. 2242–2249, IEEE, 2012.
- [5] R. Ouyang, K. H. Low, J. Chen, and P. Jaillet, "Multi-robot active sensing of non-stationary Gaussian process-based environmental phenomena," in *AAMAS*, 2014.
- [6] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in gaussian processes," in *ICML*, 2005.
- [7] J. Chen, K. H. Low, Y. Yao, and P. Jaillet, "Gaussian process decentralized data fusion and active sensing for spatiotemporal traffic modeling and prediction in mobility-on-demand systems," *IEEE ASE*, 2015.
- [8] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *Journal of Machine Learning Research*, 2008.
- [9] S. Särkkä and J. Hartikainen, "Infinite-dimensional kalman filtering approach to spatio-temporal gaussian process regression," in *International Conference on Artificial Intelligence and Statistics*, pp. 993–1001, 2012.
- [10] R. Ramathitima, M. Whitzer, S. Bhattacharya, and V. Kumar, "Automated creation of topological maps in unknown environments using a swarm of resource-constrained robots," *IEEE Robotics and Automation Letters*, 2016.
- [11] D. Palossi, M. Furci, R. Naldi, A. Marongiu, L. Marconi, and L. Benini, "An energy-efficient parallel algorithm for real-time near-optimal uav path planning," in *Proceedings of the ACM International Conference on Computing Frontiers*, ACM, 2016.
- [12] M. d. L. Braga, A. d. J. dos Santos, A. C. Pedroza, L. H. M. Costa, M. D. de Amorim, and Y. Ghamri-Doudane, "Anytime route planning with constrained devices," *Computers & Electrical Engineering*, 2016.
- [13] S. K. Gan, R. Fitch, and S. Sukkarieh, "Online decentralized information gathering with spatial-temporal constraints," *Autonomous Robots*, vol. 37, no. 1, pp. 1–25, 2014.
- [14] K. Yang, S. K. Gan, and S. Sukkarieh, "Information-theoretic exploration of unmanned aerial vehicle in unknown cluttered environment," in *Handbook of Unmanned Aerial Vehicles*, 2015.
- [15] N. Cao, K. H. Low, and J. 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*, International Foundation for Autonomous Agents and Multiagent Systems, 2013.
- [16] J. Yu, M. Schwager, and D. Rus, "Correlated orienteering problem and its application to persistent monitoring tasks," *IEEE Transactions on Robotics*, vol. 32, no. 5, pp. 1106–1118, 2016.
- [17] Z. Ma, K. Yin, L. Liu, and G. S. Sukhatme, "A spatio-temporal representation for the orienteering problem with time-varying profits," *arXiv preprint arXiv:1611.08037*, 2016.
- [18] A. Konak, D. W. Coit, and A. E. Smith, "Multi-objective optimization using genetic algorithms: A tutorial," *Reliability Engineering & System Safety*, vol. 91, no. 9, pp. 992–1007, 2006.
- [19] A. Viseras, T. Wiedemann, C. Manss, L. Magel, J. Mueller, D. Shutin, and L. Merino, "Decentralized multi-agent exploration with online-learning of gaussian processes," in *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pp. 4222–4229, IEEE, 2016.
- [20] C. Stachniss, G. Grisetti, and W. Burgard, "Information gain-based exploration using rao-blackwellized particle filters," in *Robotics: Science and Systems*, 2005.
- [21] N. D. Lawrence, "Gaussian process latent variable models for visualisation of high dimensional data," *Advances in neural information processing systems*, 2004.