

Multi-Robot Active Sensing of Non-Stationary Gaussian Process-Based Environmental Phenomena

Ruofei Ouyang[†], Kian Hsiang Low[†], Jie Chen[§], and Patrick Jaillet[‡]

Department of Computer Science, National University of Singapore, Republic of Singapore[†]

Singapore-MIT Alliance for Research and Technology, Republic of Singapore[§]

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology[‡]

{ouyang, lowkh}@comp.nus.edu.sg[†], chenjie@smart.mit.edu[§], jaillet@mit.edu[‡]

ABSTRACT

A key challenge of environmental sensing and monitoring is that of sensing, modeling, and predicting large-scale, spatially correlated environmental phenomena, especially when they are unknown and non-stationary. This paper presents a *decentralized multi-robot active sensing* (DEC-MAS) algorithm that can efficiently coordinate the exploration of multiple robots to gather the most informative observations for predicting an unknown, non-stationary phenomenon. By modeling the phenomenon using a *Dirichlet process mixture of Gaussian processes* (DPM-GPs), our work here is novel in demonstrating how DPM-GPs and its structural properties can be exploited to (a) formalize an active sensing criterion that trades off between gathering the most informative observations for estimating the unknown, non-stationary spatial correlation structure vs. that for predicting the phenomenon given the current, imprecise estimate of the correlation structure, and (b) support efficient decentralized coordination. We also provide a theoretical performance guarantee for DEC-MAS and analyze its time complexity. We empirically demonstrate using two real-world datasets that DEC-MAS outperforms state-of-the-art MAS algorithms.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Stochastic processes; I.2.9 [Robotics]: Autonomous vehicles

General Terms

Algorithms, Performance, Experimentation, Theory

Keywords

Dirichlet process mixture of Gaussian processes; multi-robot exploration and mapping; adaptive sampling; active learning

1. INTRODUCTION

A key challenge of environmental sensing and monitoring is that of sensing, modeling, and predicting complex urban and natural environmental phenomena, which are typically characterized by spatially correlated measurements [15]. To tackle this challenge, recent research efforts in the robotics community have focused on developing *multi-robot active sensing* (MAS) algorithms: Their objective is to coordinate the exploration of a team of mobile robots to actively

Appears in: *Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), May 5-9, 2014, Paris, France.* Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: Real-world non-stationary environmental phenomena: (a) Plankton density (chl-a) phenomenon (measured in mg/m^3) in log-scale in Gulf of Mexico, and (b) traffic (road speeds) phenomenon (measured in km/h) over an urban road network.

gather the most informative observations for predicting a spatially varying phenomenon of interest while being subject to resource cost constraints (e.g., number of deployed robots, energy consumption, mission time). To achieve this, a number of MAS algorithms [1, 3, 4, 9, 11, 12, 13, 14, 20, 22] have modeled the phenomenon as a *Gaussian process* (GP) [2], which allows its spatial correlation structure to be formally characterized and its predictive uncertainty to be formally quantified (e.g., based on mean-squared error, entropy, or mutual information criterion) and subsequently exploited for directing the robots to explore its highly uncertain areas. In order not to incur high computational expense, these algorithms have assumed the spatial correlation structure to be known (or estimated crudely using sparse prior data) and stationary (i.e., degree of smoothness in the spatial variation of the measurements is the same across the entire phenomenon), properties of which are often violated in real-world environmental sensing applications and limited to small-scale phenomena [21].

In practice, the spatial correlation structure of possibly large-scale environmental phenomena is usually not known and non-stationary [21] (i.e., separate areas of a phenomenon exhibit different local degrees of smoothness in the spatial variation of the measurements; e.g., see Fig. 1). For example, in some ocean phenomena (e.g., temperature, salinity, sea surface height), their measurements far offshore are more smoothly varying (i.e., more spatially correlated) in the cross-shore direction than nearshore [10]. Urban traffic networks also display non-stationary phenomena (e.g., traffic speeds, taxi demands), which pose important considerations to traffic routing and signal control. Existing MAS algorithms can still be used for sampling a non-stationary phenomenon by assuming, albeit incorrectly, its spatial cor-

relation structure to be known and stationary in order to preserve time efficiency. So, though they can gather the most informative observations under an assumed stationary correlation structure, they will perform sub-optimally with respect to the true non-stationary correlation structure.

A more desirable MAS algorithm should instead be designed to consider the informativeness of its selected observations for both estimating the unknown spatial correlation structure of a phenomenon as well as predicting the phenomenon given the true correlation structure. According to previous geostatistical studies [16], the most informative observations that are gathered for achieving the former active sensing criterion are not necessarily as informative for satisfying the latter. This raises a fundamental issue faced by active sensing: How can a MAS algorithm trade off between these two possibly conflicting criteria? The work of [7] has addressed such a trade-off for sampling stationary phenomena only and does not tackle the issue of time efficiency with multi-robot coordination.

This paper presents a *decentralized multi-robot active sensing* (DEC-MAS) algorithm that can efficiently coordinate the exploration of multiple robots to jointly optimize the above trade-off for sampling unknown, non-stationary environmental phenomena. Our DEC-MAS algorithm models a non-stationary phenomenon as a *Dirichlet process mixture of Gaussian processes* (DPM-GPs) (Section 2): Using the gathered observations, DPM-GPs can learn to automatically partition the phenomenon into separate local areas, each of which comprises measurements that vary according to a stationary spatial correlation structure and can thus be modeled by a locally stationary GP. The main contributions of our work here are novel in demonstrating how DPM-GPs and its structural properties can be exploited to (a) formalize an active sensing criterion that trades off between gathering the most informative observations for estimating the unknown partition (i.e., a key component of the non-stationary correlation structure) vs. that for predicting the phenomenon given the current, possibly imprecise estimate of the partition (Section 3), and (b) support effective and efficient decentralized coordination (Section 4). We also provide a theoretical performance guarantee for DEC-MAS and analyze its time complexity. Finally, we empirically demonstrate using two real-world datasets that DEC-MAS outperforms the state-of-the-art MAS algorithms (Section 5).

2. MODELING A PHENOMENON

2.1 Gaussian Process (GP)

A GP [2] can be used to model a spatially varying phenomenon as follows: The phenomenon is defined to vary as a realization of a GP. Let V denote a set of sampling units representing the domain of the phenomenon such that each sampling unit $x \in V$ is specified by a d -dimensional feature vector and is also associated with a realized (random) measurement y_x (Y_x if x is sampled/observed (unobserved)). Let $\{Y_x\}_{x \in V}$ denote a GP, that is, every finite subset of $\{Y_x\}_{x \in V}$ has a multivariate Gaussian distribution [2, 19]. The GP is fully specified by its *prior* mean $\mu_x \triangleq \mathbb{E}[Y_x]$ and covariance $\sigma_{xx'}|\theta \triangleq \text{cov}[Y_x, Y_{x'}|\theta]$ for all $x, x' \in V$, the latter of which characterizes the spatial correlation structure of the phenomenon and can be defined using a covariance function parameterized by θ , as described later.

Supposing a column vector y_D of realized measurements

is available for some set $D \subset V$ of observed sampling units, the GP can exploit these observations to predict the measurements for any set $X \subseteq V \setminus D$ of unobserved sampling units as well as provide their corresponding predictive uncertainties using the following Gaussian *posterior* mean vector and covariance matrix, respectively:

$$\mu_{X|D,\theta} \triangleq \mu_X + \Sigma_{XD|\theta} \Sigma_{DD|\theta}^{-1} (y_D - \mu_D) \quad (1)$$

$$\Sigma_{XX|D,\theta} \triangleq \Sigma_{XX|\theta} - \Sigma_{XD|\theta} \Sigma_{DD|\theta}^{-1} \Sigma_{DX|\theta} \quad (2)$$

where μ_X (μ_D) is a column vector with mean components μ_x for all $x \in X$ ($x \in D$), $\Sigma_{XX|\theta}$ ($\Sigma_{DD|\theta}$) is a covariance matrix with covariance components $\sigma_{xx'}|\theta$ for all $x, x' \in X$ ($x, x' \in D$), $\Sigma_{XD|\theta}$ is a covariance matrix with covariance components $\sigma_{xx'}|\theta$ for all $x \in X, x' \in D$, and $\Sigma_{DX|\theta}$ is the transpose of $\Sigma_{XD|\theta}$. The posterior covariance matrix $\Sigma_{XX|D,\theta}$ (2), which is independent of the measurements y_D , can be used to quantify the uncertainty of the predictions through, for example, the Gaussian posterior joint entropy:

$$\mathbb{H}[Y_X|y_D, \theta] \triangleq \frac{1}{2} \log(2\pi e)^{|X|} |\Sigma_{XX|D,\theta}|. \quad (3)$$

We will focus on using this entropy-based measure of predictive uncertainty in this paper.

A GP can model a stationary phenomenon by defining its prior covariance $\sigma_{xx'}|\theta$ using a stationary covariance function [19], that is, it is a function of $x - x'$. Hence, it is invariant to translations in the domain V . A common choice is the squared exponential covariance function:

$$\sigma_{xx'}|\theta \triangleq \sigma_s^2 \exp\left(-\frac{1}{2} \sum_{i=1}^d \left(\frac{x_i - x'_i}{\ell_i}\right)^2\right) + \sigma_n^2 \delta_{xx'} \quad (4)$$

where x_i (x'_i) is the i -th component of the d -dimensional feature vector x (x'), the set of hyperparameters $\theta \triangleq \{\sigma_s^2, \sigma_n^2, \ell_1, \dots, \ell_d\}$ are, respectively, signal and noise variances and length-scales, and $\delta_{xx'}$ is a Kronecker delta that is 1 if $x = x'$ and 0 otherwise. Intuitively, the signal and noise variances describe, respectively, the intensity and noise of the measurements while each length-scale ℓ_i controls the degree of smoothness in the spatial variation of the measurements (i.e., spatial correlation or “similarity” between measurements) with respect to the i -th feature component. If the hyperparameters are not known, they can be trained using the available observations via *maximum likelihood estimation* (MLE) [19], that is, by choosing θ that maximizes the log marginal likelihood $\log p(y_D|\theta) =$

$$-\frac{1}{2} (y_D - \mu_D)^\top \Sigma_{DD|\theta}^{-1} (y_D - \mu_D) - \frac{1}{2} \log(2\pi)^{|D|} |\Sigma_{DD|\theta}|. \quad (5)$$

Similarly, a GP can model a non-stationary phenomenon by specifying its prior covariance with a non-stationary covariance function, the choice of which involves a trade-off between the richness of the resulting GP model vs. computational efficiency. For example, the simple non-stationary polynomial and neural network covariance functions [19] only need a few hyperparameters to be determined. But, they do not exhibit a desirable *locality property*¹ that holds for many stationary covariance functions (e.g., (4)) and, more importantly, has been widely exploited by existing MAS algorithms mentioned in Section 1 to achieve time efficiency. On the other hand, the complex non-stationary version of

¹The locality property [8] states that the spatial correlation of measurements between sampling units decreases to zero with increasing distance between them.

Matérn covariance function [17] requires a large number of hyperparameters to be specified. Though it can capture the locality property, the training of its hyperparameters, when unknown, is computationally expensive. An alternative to using a single GP is to consider modeling the non-stationary phenomenon with a mixture of GPs that can provide a fine balance between richness and efficiency as well as a useful structural property to be exploited by our DEC-MAS algorithm, as described next.

2.2 Dirichlet Process Mixture of Gaussian Processes (DPM-GPs)

It is often observed (e.g., see Fig. 1) that the measurements in separate areas of a non-stationary phenomenon vary according to different locally stationary spatial correlation structures [21]. Such a phenomenon can be modeled with high fidelity by a Dirichlet process mixture of locally stationary GPs [18], which offers the following representational and computational advantages over a single non-stationary GP (Section 2.1): (a) It preserves the use of the well-studied and widely-applied stationary covariance functions, many of which exhibit the locality property (Section 2.1) and are computationally friendly with only a few (unknown) hyperparameters to be trained, (b) the required number of locally stationary GPs can automatically grow with the increasing complexity of the phenomenon, and (c) each locally stationary GP only incurs cubic time in the size of the observations that are local to its corresponding area of prediction instead of over the entire phenomenon.

A DPM-GPs can model a non-stationary phenomenon as follows: The phenomenon is defined to vary as a realization of a DPM-GPs. Let its number of locally stationary GP components be denoted by K . For each GP component $k = 1, \dots, K$, its prior covariance characterizes a locally stationary spatial correlation structure and is defined using a stationary covariance function parameterized by θ_k . In order to estimate the unknown θ_k using MLE (5), the measurements y_{D_k} (where $D_k \subseteq D$) that are induced by GP component k have to be identified first. That is, every observed sampling unit $x \in D$ has to be associated with a realized component label denoted by z_x and $D_k \triangleq \{x \in D | z_x = k\}$. To realize these component labels $z_D \triangleq \{z_x\}_{x \in D}$, we use Gibbs sampling, as detailed next.

Each random component label, denoted by Z_x , for all $x \in D$ follows a sampling unit-dependent Dirichlet process prior:

$$p(Z_x = k | z_{D \setminus \{x\}}, \theta_k) = \begin{cases} \frac{n_{xk}}{|D| - 1 + \alpha} & \text{if } k \leq K, \\ \frac{\alpha}{|D| - 1 + \alpha} & \text{if } k = K + 1, \end{cases} \quad (6)$$

where $n_{xk} \triangleq (|D| - 1) (\sum_{x' \in N_x} \sigma_{xx'|\theta_k} \delta_{z_{x'}k}) / (\sum_{x' \in N_x} \sigma_{xx'|\theta_k})$, $N_x \triangleq \{x' \in D \setminus \{x\} | d_G(x, x') \leq \gamma\}$ for some $\gamma > 0$, $d_G(x, x')$ is the shortest path length between sampling units x and x' with respect to the topology of a graph G induced from V to be traversed by the robots (Section 3), $\sigma_{xx'|\theta_k}$ is previously defined in (4), and α denotes a concentration parameter. The Dirichlet process prior (6) can be understood as follows: When $k = 1, \dots, K$, the probability of the observation at x being induced by GP component k is proportional to the number of neighboring observed sampling units with the same component label k weighted by their proximity to x . Its probability of being induced by a new GP component $K + 1$ is proportional to α . Hence, α controls the addition of new GP components. For the new GP component $K + 1$,

its θ_{K+1} is sampled from a pre-defined uniform distribution.

Given the realized measurements y_D for the set D of observed sampling units, the Dirichlet process prior can be updated using Bayes' rule to the following posterior:

$$p(Z_x = k | z_{D \setminus \{x\}}, y_D, \theta_k) \propto \begin{cases} p(y_x | Z_x = k, y_{D_k \setminus \{x\}}, \theta_k) p(Z_x = k | z_{D \setminus \{x\}}, \theta_k) & \text{if } k \leq K, \\ p(y_x | Z_x = k, \theta_k) p(Z_x = k | z_{D \setminus \{x\}}, \theta_k) & \text{if } k = K + 1, \end{cases} \quad (7)$$

where

$p(y_x | Z_x = k, y_{D_k \setminus \{x\}}, \theta_k) \sim \mathcal{N}(\mu_{x|D_k \setminus \{x\}, \theta_k}, \Sigma_{xx|D_k \setminus \{x\}, \theta_k})$ for $k \leq K$ and $p(y_x | Z_x = K + 1, \theta_{K+1}) \sim \mathcal{N}(\mu_x, \sigma_{xx|\theta_{K+1}})$. It can be seen from $p(y_x | Z_x = k, y_{D_k \setminus \{x\}}, \theta_k)$ that an observation induced by a GP component is conditionally independent of the observations induced by the other GP components, a structural property of which will be exploited by our DEC-MAS algorithm (Sections 3 and 4).

Using the posterior (7), Gibbs sampling [5] is performed (starting with $K = 1$) to realize the component labels z_D . Given z_D , θ_k can now be trained using MLE (5). Such a process of Gibbs sampling followed by MLE is iterated until the values of z_D stabilize or a user-defined limit is reached.

Given the realized measurements y_D and component labels z_D for observed sampling units D , the DPM-GPs can exploit them to predict the measurement for an unobserved sampling unit x by aggregating the predictions of the K GP components weighted by their probability of inducing it:

$$\mu_{x|D, \theta} = \sum_{k=1}^K \mu_{x|D_k, \theta_k} p(Z_x = k | z_D, \theta_k) \quad (8)$$

where $\theta \triangleq \{\theta_1, \dots, \theta_K\}$ and $p(Z_x = k | z_D, \theta_k)$, which is defined in a similar way to (6), can be used to estimate the unknown partition of the phenomenon.

3. MULTI-ROBOT ACTIVE SENSING (MAS)

Define a directed graph $G \triangleq (V, E)$ where the domain V of a phenomenon is connected by a set $E \subseteq V \times V$ of edges such that there is an edge (x, x') if and only if a robot can traverse from $x \in V$ to $x' \in V$ within some user-defined cost constraint (e.g., time interval, traveling distance). The MAS problem is then formulated as follows: Supposing the robots have previously observed the measurements y_D from a set $D \subset V$ of sampling units and used these observations to estimate their corresponding component labels z_D by Gibbs sampling and the hyperparameters θ of the DPM-GPs by MLE (Section 2.2), they have to coordinate to jointly select the next most informative set X^* of sampling units (i.e., with corresponding measurements and component labels of maximum joint entropy) to be observed:

$$X^* = \arg \max_X \mathbb{H}[Y_X, Z_X | y_D, z_D, \theta]. \quad (9)$$

The next possible sampling unit to be observed by each robot is constrained to be selected from one that is adjacent to the robot's current residing sampling unit in G . Using chain rule for entropy, it can be shown that these maximum entropy sampling units X^* minimize the posterior joint entropy (i.e., $\mathbb{H}[Y_{V \setminus (D \cup X^*)}, Z_{V \setminus (D \cup X^*)} | Y_{X^*}, Z_{X^*}, y_D, z_D, \theta]$) of the measurements and component labels for the remaining unobserved sampling units (i.e., $V \setminus (D \cup X^*)$) in the phenomenon. $\mathbb{H}[Y_{V \setminus (D \cup X^*)}, Z_{V \setminus (D \cup X^*)} | Y_{X^*}, Z_{X^*}, y_D, z_D, \theta] = \mathbb{H}[Z_{V \setminus (D \cup X^*)} | Z_{X^*}, z_D, \theta] + \mathbb{H}[Y_{V \setminus (D \cup X^*)} | Y_{X^*}, Z_{V \setminus (D \cup X^*)}, y_D, z_D, \theta]$ by chain rule for entropy. So, the choice of X^* (9) jointly optimizes a trade-off between gathering the most informative

$$\begin{aligned}
& \max_X \sum_{x \in X} \mathbb{H}[Z_x | z_D, \theta] + \sum_{k=1}^K \widehat{\mathbb{H}}[Y_{X_k} | \widehat{z}_{X_k}, y_{D_k}, \theta_k] \\
&= \max_{\bigcup_{n=1}^N X_n} \sum_{n=1}^N \sum_{x \in X_n} \mathbb{H}[Z_x | z_D, \theta] + \sum_{k=1}^K \mathbb{H}[Y_{X_{k_n}} | \widehat{z}_{X_{k_n}}, y_{D_k}, \theta_k] \\
&= \sum_{n=1}^N \max_{X_n} \sum_{x \in X_n} \mathbb{H}[Z_x | z_D, \theta] + \sum_{k=1}^K \mathbb{H}[Y_{X_{k_n}} | \widehat{z}_{X_{k_n}}, y_{D_k}, \theta_k]
\end{aligned} \tag{16}$$

where the first equality is due to (15). More importantly, the last equality can be solved in a partially decentralized manner by each disjoint subset \mathcal{V}_n of robots for $n = 1, \dots, N$:

$$\widehat{X}_n = \arg \max_{X_n} \sum_{x \in X_n} \mathbb{H}[Z_x | z_D, \theta] + \sum_{k=1}^K \mathbb{H}[Y_{X_{k_n}} | \widehat{z}_{X_{k_n}}, y_{D_k}, \theta_k]. \tag{17}$$

The degree of decentralization for our DEC-MAS algorithm (17) can be varied by controlling ε : Increasing ε causes more robots to become isolated vertices in \mathcal{G} , thus decreasing the size $\eta \triangleq \max_n |\mathcal{V}_n|$ of its largest connected component and entailing higher degree of decentralization.

Let

$$\xi \triangleq \max_{k,n, X_{k_n}, i, i'} \left| \left[\Sigma_{X_{k_n} X_{k_n} | D_k, \theta_k}^{-1} \right]_{ii'} \right| \tag{18}$$

and $\epsilon \triangleq 0.5K \log 1 / \left(1 - (|\mathcal{V}|^{1.5} \eta \xi \varepsilon)^2 \right)$. We prove in the theoretical result below that $\widehat{X} = \bigcup_{n=1}^N \widehat{X}_n$ is guaranteed to achieve an entropy $\widehat{\mathbb{H}}[Y_{\widehat{X}}, Z_{\widehat{X}} | y_D, z_D, \theta]$ (i.e., by plugging \widehat{X} into (12)) that is at most ϵ less than the maximum entropy $\widetilde{\mathbb{H}}[Y_{\widetilde{X}}, Z_{\widetilde{X}} | y_D, z_D, \theta]$ achieved by \widetilde{X} (11):

THEOREM 1 (PERFORMANCE GUARANTEE). *If $|\mathcal{V}|^{1.5} \eta \xi \varepsilon < 1$, then $\widehat{\mathbb{H}}[Y_{\widehat{X}}, Z_{\widehat{X}} | y_D, z_D, \theta] - \widetilde{\mathbb{H}}[Y_{\widetilde{X}}, Z_{\widetilde{X}} | y_D, z_D, \theta] \leq \epsilon$.*

The proof of Theorem 1 is given in Appendix A. The implication of Theorem 1 is that our DEC-MAS algorithm (17) is competitive (i.e., small ϵ) as compared to the CEN-MAS algorithm (11) when (a) the number $|\mathcal{V}|$ of robots is not too large, (b) the largest connected component of η robots being formed in \mathcal{G} is reasonably small, (c) the minimum required correlation ε between the next possible sampling units to be observed by adjacent robots is kept low, and (d) the number K of GP components is small.

4.1 Time and Communication Complexity

In this subsection, we will analyze the time and communication complexity of our DEC-MAS algorithm. Suppose that the observations are distributed evenly among the K GP components and denote the maximum out-degree and in-degree of G by δ and δ' , respectively. Then, $|N_x| \leq \Delta \triangleq (\delta + \delta')^\gamma$ for all $x \in D$. Gibbs sampling for estimating the component labels z_D followed by MLE for estimating the hyperparameters θ (Section 2.2) incur $\mathcal{O}(M|D|K((|D|/K)^3 + \Delta))$ time over M iterations. Our DEC-MAS algorithm (17) incurs $\mathcal{O}(K(|D|/K)^3 + \eta \delta^\gamma (K\Delta + (|D|/K)^2 + \eta^2))$ time. By setting $\eta = |\mathcal{V}|$, it yields the time complexity of the CEN-MAS algorithm (11) for comparison.

Central to the efficiency of our DEC-MAS algorithm is the requirement of a small η (i.e., size of largest connected component of robots being formed in \mathcal{G} to coordinate their active sensing), which is in fact achieved in practice, as explained by the following observations: For a GP component with small spatial correlation, the posterior entropy of the measurements in the unobserved part of its local area of prediction remains high after sampling, hence attracting more

robots to explore it. But, its small spatial correlation entails high degree of decentralization (13), thus resulting in a small η . On the other hand, for a GP component with large spatial correlation, the posterior entropy of the measurements in the unobserved part of its local area of prediction becomes low after sampling, hence attracting fewer robots to explore it. So, a small η is also maintained.

For our DEC-MAS algorithm, each robot broadcasts $\mathcal{O}(|\mathcal{V}|)$ -sized and $\mathcal{O}(1)$ -sized messages on its adjacency information and new observation, respectively.

5. EXPERIMENTS AND DISCUSSION

This section evaluates the active sensing performance and time efficiency of DEC-MAS empirically on two real-world datasets featuring non-stationary phenomena: (a) June 2012 MODIS plankton density (chlorophyll-a) data of Gulf of Mexico (Fig. 1a) discretized into a 60×60 grid of sampling locations/units and bounded within lat. $28.175 - 29.975$ N and lon. $87.675 - 89.475$ W. The mean density is 4.5 mg/m^3 and standard deviation is 9.8 mg/m^3 ; (b) Traffic speeds data along 775 road segments (including highways, arterials, slip roads, etc.) of an urban road network (Fig. 1b) during the evening peak hours on April 20, 2011. The mean speed is 52.8 km/h and standard deviation is 21.1 km/h . Each sampling unit (i.e., road segment) is specified by a 4-dimensional feature vector: length, number of lanes, speed limit, and direction. This non-stationary traffic phenomenon is modeled using a Dirichlet process mixture of stationary relational GPs; the relational GP is previously developed in [4] and its stationary correlation structure can exploit both the road segment features and road network topology information.

For each dataset, 5% of the data are randomly selected as prior observations to estimate their corresponding prior component labels z_D by Gibbs sampling and the prior hyperparameters θ of the DPM-GPs by MLE (Section 2.2). Subsequently, they are constantly updated using the new observations gathered by running DEC-MAS repeatedly. For DEC-MAS, ε (13) is set to 0.1. The experiments are run on a PC with Intel® Core™2 Quad CPU Q9550 at 2.83 GHz. The results shown below are averaged over 40 trials of randomly selected initial robots' residing sampling units.

Performance metrics. The first metric evaluates active sensing performance of a tested MAS algorithm: It measures *root mean squared error* (RMSE) $\sqrt{\sum_{x \in V} (\mu_{x|D, \theta} - y_x)^2 / |V|}$ over domain V of the phenomenon that results from using the posterior mean $\mu_{x|D, \theta}$ of the algorithm's utilized model (i.e., (1) of GP or (8) of DPM-GPs with stationary covariance function (4)) to predict the measurements for the remaining unobserved sampling units $V \setminus D$ given the gathered observations. The second metric evaluates the time efficiency and scalability of a tested MAS algorithm by measuring its incurred time.

Comparison of MAS algorithms. The performance of our DEC-MAS algorithm is compared to that of the state-of-the-art MAS algorithms, as listed in Table 1 and briefly described next: The *centralized maximum entropy sampling* (CEN-MES) algorithm [13] repeatedly selects the next set X of sampling units to be observed that maximizes (3) based on a stationary GP model. After gathering the observations, CEN-MES can alternatively use DPM-GPs (instead of GP) for prediction (i.e., (8)) and we call this CEN-MES+D. The *partially decentralized maximum entropy sampling* (DEC-

Algorithm	Model	Criterion
CEN-MES [13]	GP	(3)
DEC-MES [4]	GP	(3)
CEN-MES+D	GP (active sensing) DPM-GPs (prediction)	(3)
MAX-SUM [20]	DPM-GPs	(11)
CEN-MAS	DPM-GPs	(11)
DEC-MAS	DPM-GPs	(17)
DEC-MAS-C	DPM-GPs	(19)

Table 1: Comparison of MAS algorithms (Each algorithm exploits a single model for both active sensing and prediction, except for CEN-MES+D).

MES) algorithm [4] exploits a similar notion of the coordination graph to split a robot team into disjoint sub-teams, each of which runs CEN-MAS separately without coordinating with other sub-teams. The MAX-SUM algorithm [20] is a general-purpose iterative solver for distributed constraint optimization problems. In [20], MAX-SUM is only used to optimize (3) based on the GP model; it does not utilize DPM-GPs nor optimize our novel MAS criterion (11), which are done here. Unlike DEC-MAS, the performance guarantee of MAX-SUM offers a non-informative, loose worst-case approximation ratio that only holds for tree-like coordination structures. Lastly, to show the importance of observing sampling units with highly uncertain component labels, the first summation term in (17) is removed to yield

$$\max_{X_n} \sum_{k=1}^K \mathbb{H}[Y_{X_{kn}} | \hat{z}_{X_{kn}}, y_{D_k}, \theta_k], \quad (19)$$

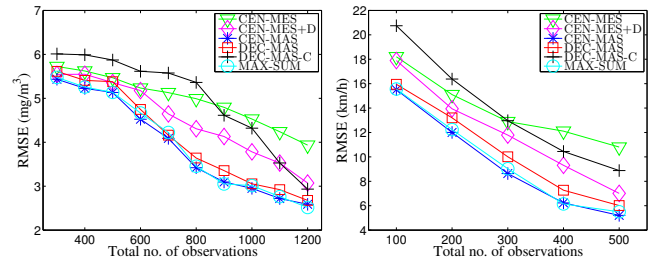
which we call DEC-MAS-C. Note that it is prohibitively expensive to compare with the maximum mutual information-based algorithm of [8], which scales poorly with increasing domain size $|V|$ and is hence not practical for real-time active sensing. For example, it incurred > 62 hours to generate paths for 3 robots to sample a total of 267 observations in a grid of $|V| = 1424$ sampling units, as reported in [14].

5.1 Results and Analysis

A. Effect of criterion on predictive performance. Fig. 2 shows results of the predictive performance using varying number $|D|$ of observations gathered by $|\mathcal{V}| = 4$ robots running the tested algorithms. The observations are as follows:

(I) The algorithms optimizing active sensing criterion (11) or (17) based on DPM-GPs (i.e., CEN-MAS, DEC-MAS, and MAX-SUM) can achieve the best predictive performance (i.e., lowest RMSE) due to the following reasons: (a) DPM-GPs can model and predict the non-stationary phenomena better than a stationary GP, as observed in the performance improvement of CEN-MES+D over CEN-MES by using DPM-GPs (instead of GP) for prediction, and (b) algorithms optimizing the criteria (11) or (17) can gather more informative observations than algorithms using criterion (3), as observed in the performance improvement of CEN-MAS, DEC-MAS, and MAX-SUM over CEN-MES+D while using DPM-GPs for prediction.


(II) More superior predictive performance can be achieved by jointly optimizing the trade-off between observing sampling units with most uncertain component labels vs. that with most uncertain measurements given the current, imprecise estimate of their labels than by solely addressing the latter criterion; this is observed in the more superior performance of CEN-MAS, DEC-MAS, and MAX-SUM over



(a) Plankton density

(b) Traffic speeds

Figure 2: Graphs of predictive performance vs. total no. $|D|$ of observations gathered by $|\mathcal{V}| = 4$ robots.



(a) Plankton density

(b) Traffic speeds

Figure 3: Graphs of predictive performance vs. no. $|\mathcal{V}|$ of deployed robots gathering a total of (a) $|D| = 1200$ and (b) $|D| = 500$ observations from plankton density and traffic phenomena, respectively.

DEC-MAS-C, the latter of which neglects observing sampling units with highly uncertain labels (19).

(III) DEC-MAS optimizing criterion (17) can achieve predictive performance close to that of CEN-MAS and MAX-SUM using criterion (11). It is prohibitively expensive to obtain results for CEN-MAS with $|\mathcal{V}| > 4$ robots. So, we will only present results for decentralized algorithms from now on.

B. Effect of decentralization on predictive performance. Fig. 3 shows results of the predictive performance using a total of $|D| = 1200$ and $|D| = 500$ observations gathered from plankton density and traffic phenomena, respectively, by varying number $|\mathcal{V}|$ of robots running the tested decentralized algorithms. The observations are as follows:

(I) The predictive performance of all decentralized algorithms improve with increasing number of robots because every robot is tasked to gather less observations and their performance are thus less adversely affected by their greedy selection of maximum-entropy sampling units. Consequently, more informative unobserved sampling units are explored.

(II) DEC-MAS performs significantly better than DEC-MES and DEC-MAS-C due to the same reasons as that given in the previous observations A(I) and A(II), respectively.

(III) DEC-MAS can achieve predictive performance comparable to that of MAX-SUM. Intuitively, MAX-SUM exploits and exchanges additional coordination information between robots in different connected components formed by DEC-MAS, but this results in little performance improvement of MAX-SUM over DEC-MAS. We will also see later that MAX-SUM is less computationally efficient and significantly less scalable than DEC-MAS in the number of robots.

C. Effect of decentralization on time efficiency and scalability. Fig. 4 shows results of the incurred time of the tested algorithms with varying number of observations and robots. The observations are as follows:

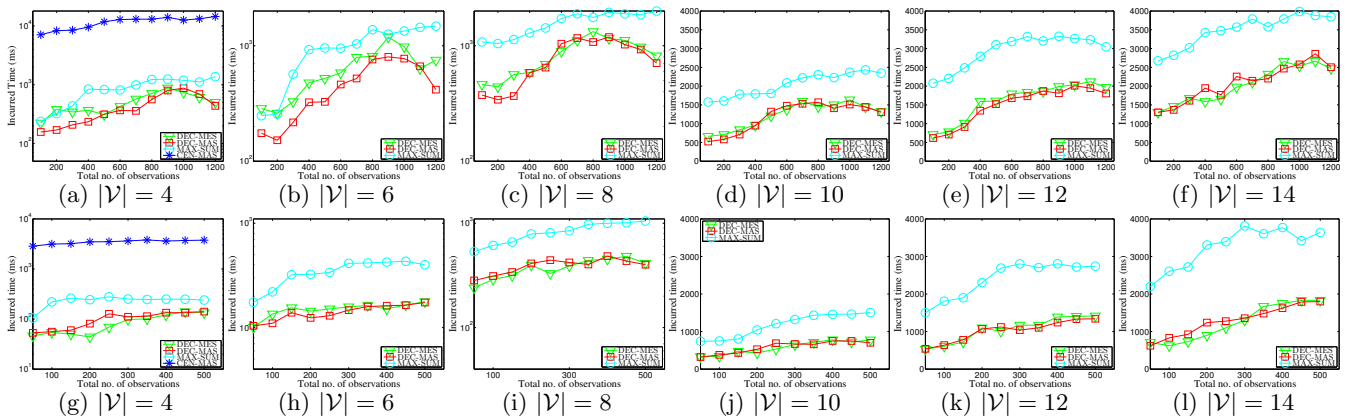


Figure 4: Graphs of incurred time vs. total no. $|D|$ of observations gathered from (a-f) plankton density and (g-l) traffic phenomena by varying no. $|V|$ of robots.

(I) CEN-MAS incurs at least 1 order of magnitude more time than the decentralized algorithms for $|V| = 4$ robots.

(II) DEC-MAS incurs computational time less than or comparable to that of DEC-MES: Even though DEC-MAS incurs additional time needed to estimate the component labels and compute the entropy of labels in (17), it saves time in the following aspects: (a) As mentioned previously in Section 2.2, DPM-GPs (i.e., used by DEC-MAS) offers the computational advantage over a single GP (i.e., used by DEC-MES) that each GP component only incurs cubic time in the size of the observations that are local to its corresponding area of prediction instead of over the entire phenomenon; and (b) DEC-MAS tends to form smaller connected components than DEC-MES due to the structural property of DPM-GPs that requires two robots to coordinate their active sensing only when some pair of their next possible sampling units to be observed are associated with the same GP component (Section 4), and also due to its behavior of keeping the size η of the largest connected component small, as explained in Section 4.1.

(III) DEC-MAS is more time-efficient and significantly more scalable than MAX-SUM in the number of robots (Fig. 4) while achieving comparable predictive performance (Fig. 3). MAX-SUM is computationally more expensive because it has to process the additional coordination information between robots in different connected components formed by DEC-MAS that results in little performance improvement over DEC-MAS.

6. CONCLUSION

This paper describes a novel DEC-MAS algorithm that can efficiently coordinate multiple robots in a partially decentralized manner to gather the most informative observations for predicting an unknown, non-stationary phenomenon. In particular, we demonstrate how its efficient decentralized coordination and theoretical performance guarantee can be realized by exploiting the structural property of DPM-GPs and the locality property of each stationary GP component. Empirical evaluation on two real-world datasets featuring non-stationary phenomena shows that (a) more superior active sensing performance can be achieved by optimizing our proposed MAS criterion (11) or (17) that trades off between observing sampling units with most uncertain component labels vs. that with most uncertain measurements given the current, imprecise estimate of their labels, and (b) DEC-

MAS outperforms the decentralized MAX-SUM [20] (DEC-MES [4]) algorithm in time efficiency and scalability (active sensing) while achieving comparable active sensing performance (time efficiency).

Acknowledgments. This work was supported by Singapore-MIT Alliance Research & Technology Subaward Agreements No. 41 R-252-000-527-592 & No. 47 R-252-000-509-592.

7. REFERENCES

- [1] 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 *Proc. AAMAS*, pages 7–14, 2013.
- [2] J. Chen, N. Cao, K. H. Low, R. Ouyang, C. K.-Y. Tan, and P. Jaillet. Parallel Gaussian process regression with low-rank covariance matrix approximations. In *Proc. UAI*, pages 152–161, 2013.
- [3] J. Chen, K. H. Low, and C. K.-Y. Tan. Gaussian process-based decentralized data fusion and active sensing for mobility-on-demand system. In *Proc. RSS*, 2013.
- [4] J. Chen, K. H. Low, C. K.-Y. Tan, A. Oran, P. Jaillet, J. M. Dolan, and G. S. Sukhatme. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In *Proc. UAI*, pages 163–173, 2012.
- [5] W. R. Gilks, S. Richardson, and D. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC, 1995.
- [6] I. C. F. Ipsen and D. J. Lee. Determinant approximations. Technical Report CRSC-TR03-30, Center for Research in Scientific Computation, North Carolina State University, Raleigh, NC, 2003.
- [7] A. Krause and C. Guestrin. Nonmyopic active learning of Gaussian processes: An exploration-exploitation approach. In *Proc. ICML*, pages 449–456, 2007.
- [8] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *JMLR*, 9:235–284, 2008.
- [9] N. E. Leonard, D. Paley, F. Lekien, R. Sepulchre, D. M. Fratantoni, and R. Davis. Collective motion, sensor networks and ocean sampling. *Proc. IEEE*, 95(1):48–74, 2007.

- [10] Z. Li, Y. Chao, J. C. McWilliams, and K. Ide. A three-dimensional variational data assimilation scheme for the Regional Ocean Modeling System: Implementation and basic experiments. *J. Geophys. Res.*, 113:C05002, 2008.
- [11] K. H. Low, J. Chen, J. M. Dolan, S. Chien, and D. R. Thompson. Decentralized active robotic exploration and mapping for probabilistic field classification in environmental sensing. In *Proc. AAMAS*, 2012.
- [12] K. H. Low, J. M. Dolan, and P. Khosla. Adaptive multi-robot wide-area exploration and mapping. In *Proc. AAMAS*, pages 23–30, 2008.
- [13] K. H. Low, J. M. Dolan, and P. Khosla. Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing. In *Proc. ICAPS*, pages 233–240, 2009.
- [14] K. H. Low, J. M. Dolan, and P. Khosla. Active Markov information-theoretic path planning for robotic environmental sensing. In *Proc. AAMAS*, pages 753–760, 2011.
- [15] K. H. Low, G. J. Gordon, J. M. Dolan, and P. Khosla. Adaptive sampling for multi-robot wide-area exploration. In *Proc. IEEE ICRA*, 2007.
- [16] W. G. Müller. *Collecting Spatial Data: Optimum Design of Experiments for Random Fields*. Springer, 3rd edition, 2007.
- [17] C. J. Paciorek and M. J. Schervish. Nonstationary covariance functions for Gaussian process regression. In *Proc. NIPS*, 2003.
- [18] C. E. Rasmussen and Z. Ghahramani. Infinite mixtures of Gaussian process experts. In *Proc. NIPS*, pages 881–888, 2001.
- [19] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- [20] A. Rogers, A. Farinelli, R. Stranders, and N. R. Jennings. Bounded approximate decentralised coordination via the max-sum algorithm. *AIJ*, 175(2):730–759, 2011.
- [21] P. D. Sampson, D. Damian, and P. Guttorp. Advances in modeling and inference for environmental processes with nonstationary spatial covariance. In P. Monestiez, D. Allard, and R. Froidevaux, editors, *geoENV III – Geostatistics for Environmental Applications*, volume 11 of *Quantitative Geology and Geostatistics*, pages 17–32. Springer, 2001.
- [22] A. Singh, A. Krause, C. Guestrin, and W. J. Kaiser. Efficient informative sensing using multiple robots. *JAIR*, 34:707–755, 2009.

APPENDIX

A. PROOF OF THEOREM 1

Let $\tilde{\Sigma}_{X_k X_k | D_k, \theta_k} \triangleq \Sigma_{X_k X_k | D_k, \theta_k} - \hat{\Sigma}_{X_k X_k | D_k, \theta_k}$ and ρ_k be the spectral radius of $\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k}$. We first bound ρ_k from above.

For any X_k , $\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k}$ comprises diagonal blocks of size $|X_{kn}| \times |X_{kn}|$ with components of value 0 for $n = 1, \dots, N$ and off-diagonal blocks of the form $\Sigma_{X_{kn} X_{kn'} | D_k, \theta_k}^{-1} \Sigma_{X_{kn} X_{kn'} | D_k, \theta_k}$ for $n, n' = 1, \dots, N$ and $n \neq n'$. Any pair of robots $r \in \mathcal{V}_n$ and $r' \in \mathcal{V}_{n'}$ reside in different connected

components of coordination graph \mathcal{G} and are therefore not adjacent. So, by (13),

$$\max_{i, i'} \left| \left[\Sigma_{X_{kn} X_{kn'} | D_k, \theta_k} \right]_{ii'} \right| \leq \varepsilon \quad (20)$$

for $n, n' = 1, \dots, N$ and $n \neq n'$. Using (18) and (20), each component in any off-diagonal block of $\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k}$ can be bounded as follows:

$$\max_{i, i'} \left| \left[\Sigma_{X_{kn} X_{kn'} | D_k, \theta_k}^{-1} \Sigma_{X_{kn} X_{kn'} | D_k, \theta_k} \right]_{ii'} \right| \leq |X_{kn}| \xi \varepsilon \quad (21)$$

for $n, n' = 1, \dots, N$ and $n \neq n'$. It follows from (21) that

$$\max_{i, i'} \left| \left[\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k} \right]_{ii'} \right| \leq \max_n |X_{kn}| \xi \varepsilon \leq \eta \xi \varepsilon. \quad (22)$$

The last inequality is due to $\max_n |X_{kn}| \leq \max_n |\mathcal{V}_n| \leq \eta$. Then,

$$\begin{aligned} \rho_k &\leq \left\| \tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k} \right\|_2 \\ &\leq |X_k| \max_{i, i'} \left| \left[\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}^{-1} \tilde{\Sigma}_{X_k X_k | D_k, \theta_k} \right]_{ii'} \right| \\ &\leq |\mathcal{V}| \eta \xi \varepsilon. \end{aligned} \quad (23)$$

The first two inequalities are due to standard properties of matrix norm. The last inequality follows from (22).

The rest of this proof uses the following result of [6] that is revised to reflect our notations:

THEOREM 2. *If $|X_k| \rho_k^2 < 1$, then $\log |\Sigma_{X_k X_k | D_k, \theta_k}| \leq \log |\tilde{\Sigma}_{X_k X_k | D_k, \theta_k}| \leq \log |\Sigma_{X_k X_k | D_k, \theta_k}| - \log(1 - |X_k| \rho_k^2)$ for any X_k .*

Using Theorem 2 followed by (23),

$$\begin{aligned} \log \left| \tilde{\Sigma}_{X_k X_k | D_k, \theta_k} \right| - \log \left| \Sigma_{X_k X_k | D_k, \theta_k} \right| &\leq \log \frac{1}{1 - |X_k| \rho_k^2} \\ &\leq \log \frac{1}{1 - (|\mathcal{V}|^{1.5} \eta \xi \varepsilon)^2} \end{aligned} \quad (24)$$

for any X_k .

$$\begin{aligned} &\tilde{\mathbb{H}}[Y_{\tilde{X}}, Z_{\tilde{X}} | y_D, z_D, \theta] - \tilde{\mathbb{H}}[Y_{\tilde{X}}, Z_{\tilde{X}} | y_D, z_D, \theta] \\ &= \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] - \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] \\ &\quad + \sum_{k=1} \mathbb{H}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] - \mathbb{H}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] \\ &\leq \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] - \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] \\ &\quad + \sum_{k=1} \hat{\mathbb{H}}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] - \mathbb{H}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] \\ &\leq \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] - \sum_{x \in \tilde{X}} \mathbb{H}[Z_x | z_D, \theta] \\ &\quad + \sum_{k=1} \hat{\mathbb{H}}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] - \mathbb{H}[Y_{\tilde{X}_k} | \tilde{z}_{\tilde{X}_k}, y_{D_k}, \theta_k] \\ &\leq \frac{K}{2} \log \frac{1}{1 - (|\mathcal{V}|^{1.5} \eta \xi \varepsilon)^2} \end{aligned}$$

The first equality is due to (12). The first, second, and last inequalities follow from (14) and Theorem 2, (16), and (24), respectively.