2018 IEEE Conference on Decision and Control (CDC)
Miami Beach, FL, USA, Dec. 17-19, 2018

Adaptive Sensor Selection for Monitoring Stochastic Processes

Shinkyu Park, Carlo Ratti, and Daniela Rus

***Abstract*— We investigate an adaptive sensor selection problem in which a stochastic process is monitored by multiple sensors. We design a sensor selection policy that assigns a set of sensors to collect measurements for which the sensor selection depends on previously collected measurements and auxiliary data, and is subject to a constraint on the number of sensors to be selected. We use the mutual information to assess the performance of the policy.**

The goal of this paper is to find an approximate solution to the sensor selection problem and to assess the performance of the solution. For this purpose, we define greedy adaptive policies using greedy heuristics and derive a performance bound on greedy policies with respect to the performance of adaptive policies satisfying so-called diminishing property and optimality conditions. The main result shows that the performance of a greedy adaptive policy is at least $\frac{1}{2}$ of that of the best policy satisfying these two conditions.

I. INTRODUCTION

Given a stochastic process $\{\mathbf{x}_k\}_{k=1}^N \subset \mathbb{X}$, we consider a system consisting of M sensors that monitors the process. At each time index k in $\{1, \cdots, N\}$, the system selects a set $\mathbf{S}_k$ of sensors to collect measurements $\{\mathbf{y}_k^{(j)}\}_{j \in \mathbf{S}_k} \subset \mathbb{Y}^{|\mathbf{S}_k|}$ and has access to auxiliary data $\mathbf{z}_k \in \mathbb{Z}$, where $\mathbf{y}_k^{(j)}$ and $\mathbf{z}_k$ represent noisy information on $\mathbf{x}_k$. The sensor selection at each time index k depends on previously collected sensor measurements $(\{\mathbf{y}_1^{(j)}\}_{j \in \mathbf{S}_1}, \cdots, \{\mathbf{y}_{k-1}^{(j)}\}_{j \in \mathbf{S}_{k-1}})$ and available auxiliary data $(\mathbf{z}_1, \cdots, \mathbf{z}_k)$, and is subject to the cardinality constraint $|\mathbf{S}_k| \leq s_k$ for a positive integer s_k.[1]

In this paper, we focus on finding a computationally feasible policy that governs the sensor selection process. To assess the performance of each policy μ, we adopt the mutual information (MI) between the stochastic process and measurements from selected sensors given auxiliary data:

$$\mathcal{I}(\mathbf{x}_1, \cdots, \mathbf{x}_N; \mathbf{y}_{\mathbf{S}_1^\mu}, \cdots, \mathbf{y}_{\mathbf{S}_N^\mu} \,|\, \mathbf{z}_1, \cdots, \mathbf{z}_N) \tag{1}$$

where $\mathbf{y}_{\mathbf{S}_k^\mu} \stackrel{\text{def}}{=} (\mathbf{S}_k^\mu, \{\mathbf{y}_k^{(j)}\}_{j \in \mathbf{S}_k^\mu})$ is a tuple of the sensor selection set $\mathbf{S}_k^\mu$ and collected measurements $\{\mathbf{y}_k^{(j)}\}_{j \in \mathbf{S}_k^\mu}$. For every k in $\{1, \cdots, N\}$, an adaptive policy $\mu = (\mu_1, \cdots, \mu_N)$ recursively defines the set $\mathbf{S}_k^\mu$ of sensors to collect measurements by

$$\mathbf{S}_k^\mu = \mu_k(\mathbf{y}_{\mathbf{S}_1^\mu}, \cdots, \mathbf{y}_{\mathbf{S}_{k-1}^\mu}, \mathbf{z}_1, \cdots, \mathbf{z}_k) \tag{2}$$

subject to the constraint $|\mathbf{S}_k^\mu| \leq s_k$.

Different from existing work, here we consider the adaptive setting in which the sensor selection process depends on both previously collected sensor measurements and auxiliary data. However, finding the optimal solution to optimization (1) is known to be NP-hard even for a special case of our problem [1]. To find a computationally tractable solution, we adopt greedy heuristics from [2] to define greedy adaptive policies and provide analytical results on the performance of greedy policies.

The key technical challenges lie in assessing the performance of greedy heuristics-based approaches in the adaptive sensor selection problem. In literature, under submodularity assumptions, many of notable efforts have been made to establish the performance of greedy heuristics-based approaches in sensor selection problems. However, without the assumptions, few results are known, and in many cases even checking such submodularity conditions is a non-trivial task [3]. Besides, examples in a previous study [4] suggest that MI is in general not submodular in adaptive sensor selection problems, and existing submodularity analysis cannot be extended to our problem setting.

Our paper contributes in the following ways: **(i)** We adopt greedy heuristics to define greedy adaptive policies for the sensor selection problem. **(ii)** To establish a performance bound on greedy policies, we define a class of adaptive policies satisfying diminishing property and optimality conditions[2] and show that a greedy policy performs at least $\frac{1}{2}$ of the performance of the best policy within the class. **(iii)** We establish connections between our results and those from submodular optimization literature, and show how our analysis generalizes results in the literature. **(iv)** We execute experiments in monitoring yeast population growth to demonstrate how our approach can be used in a biology application and to assess the performance of greedy policies.

We organize the paper as follows: In Section II, we proceed with introducing a motivating application of our adaptive sensor selection problem. In Section III, we provide a formal description of the adaptive sensor selection problem and the definition of greedy adaptive policies considered

The authors would like to thank Kuwait-MIT Center for Natural Resources and the Environment, NSF 1723943, Allianz, Amsterdam Institute for Advanced Metropolitan Solutions, Brose, Cisco, Ericsson, Fraunhofer Institute, Liberty Mutual Institute, Shenzhen, Singapore- MIT Alliance for Research and Technology (SMART), UBER, Victoria State Government, Volkswagen Group America, and all the members of the MIT Senseable City Lab Consortium for supporting this research.

Shinkyu Park is with Senseable City Laboratory and Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. shinkyu@mit.edu

Carlo Ratti is with Senseable City Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. ratti@mit.edu

Daniela Rus is with Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. rus@csail.mit.edu

[1]The auxiliary data $\mathbf{z}_k$ can be regarded as any information that are not subject to the cardinality constraint. See the example given in Section II.

[2]The precise definitions of the conditions are given in Section III-B.

throughout the paper. In Section IV, we provide a comparative review of relevant work in literature. In Section V, we analyze the performance of greedy adaptive policies and establish a connection to results from submodular optimization literature. In Section VI, we present experiment results on monitoring yeast population growth using the adaptive sensor selection scheme. We end the paper with conclusions and future plans in Section VII.

II. Motivating Application

An application that motivates our adaptive sensor selection problem is in monitoring human health signals from city's sewage network (see the project Underworlds [5] for more descriptions). Consider a network of M sensor-enabled sampling devices in which each device $j \in \{1, \cdots, M\}$ is capable of collecting data $\mathbf{z}_k^{(j)}$ using its on-board sensor and collecting a sewage sample from which we could extract a measurement $\mathbf{y}_k^{(j)}$ through processing the sample at biology facilities. Data collected by the network can be used to infer public health pattern in urban epidemiology studies.

One of key challenges in this application is due to the costly process of sample collection/processing. The advent of miniaturized liquid sample analysis platforms – such as miniature mass spectrometers or microfluidics devices – enables in situ detection of biological signals [6], [7]. However, measuring majority of health-related signals still requires pre-processing of the samples, which is a challenging task with in situ sensing platforms. Also comprehensive analysis of samples to obtain information on the signals at a fine-level of granularity still requires significant human involvement which is costly and there are potentially a large number of locations for sample collection. This naturally imposes budget constraints on the process such as limits on the number of samples that can be collected and processed on each experiment-day.

To address the challenge, we formulate the problem as adaptive sensor selection and find a computationally feasible solution. In the problem, as in the motivating application, we can consider that $\mathbf{y}_k^{(j)}$ is the informative measurement of $\mathbf{x}_k$ from sensor j whereas the auxiliary data $\mathbf{z}_k$ is the noisy measurement of $\mathbf{x}_k$. We present the following example to demonstrate how we can adopt the adaptive sensor selection framework for monitoring population growth. In Section VI we describe experiments using the models from the example.

***Example** II.1:* Consider M separate locations where at each location we aim at monitoring the population of a biological agent growing at the location. Let us consider that the population growth of the agent is described by a stochastic exponential growth model [8] given as follows:

$$\mathbf{x}_{k+1}^{(j)} = \mathbf{x}_k^{(j)} + r^{(j)}\mathbf{x}_k^{(j)} \left[1 - \mathbf{x}_k^{(j)}/K^{(j)}\right] + \mathbf{v}_k^{(j)} \tag{3}$$

where $\mathbf{x}_k^{(j)}$ denotes the population of the agent at location j and at time index k, and $r^{(j)}$ and $K^{(j)}$ denote the growth rate and carrying capacity, respectively. The random variables $\mathbf{x}_1^{(j)}$ and $\mathbf{v}_k^{(j)}$ describe the initial condition and noise term, respectively. Note that each $\mathbf{x}_k^{(j)}$ constitutes $\mathbf{x}_k$ as $\mathbf{x}_k = (\mathbf{x}_k^{(1)}, \cdots, \mathbf{x}_k^{(M)})$, which defines the stochastic process being monitored in our sensor selection problem.

To monitor the population growth, consider that M sensor-enabled sampling devices are deployed across all the locations where each device j is deployed at location j. For each time index k in $\{1, \cdots, N\}$, we denote $\mathbf{y}_k^{(j)}$ as a measurement obtained through collecting/processing a sample of the agent from location j, and $\mathbf{z}_k^{(j)}$ as data collected from device j using its on-board sensor. We adopt the following models to describe how $\mathbf{y}_k^{(j)}$ and $\mathbf{z}_k^{(j)}$ are related with $\mathbf{x}_k^{(j)}$:

$$\mathbf{y}_k^{(j)} = \begin{cases} \mathbf{x}_k^{(j)} + \mathbf{w}_k^{(j)} & \text{w.p. } p \\ \varnothing & \text{w.p. } 1-p \end{cases} \tag{4a}$$

$$\mathbf{z}_k^{(j)} = \mathbf{x}_k^{(j)} + \mathbf{u}_k^{(j)} \tag{4b}$$

where the random variables $\mathbf{w}_k^{(j)}$ and $\mathbf{u}_k^{(j)}$ are mutually independent disturbances for all j in $\{1, \cdots, M\}$ and k in $\{1, \cdots, N\}$, and the entropy of the random variables satisfy $h(\mathbf{w}_k^{(j)}) << h(\mathbf{u}_k^{(j)})$, i.e., the measurement $\mathbf{y}_k^{(j)}$ contains more information about $\mathbf{x}_k^{(j)}$ than does $\mathbf{z}_k^{(j)}$. We define auxiliary data $\mathbf{z}_k$ as $\mathbf{z}_k = (\mathbf{z}_k^{(1)}, \cdots, \mathbf{z}_k^{(M)})$. Given an adaptive policy $\mu = (\mu_1, \cdots, \mu_N)$, the set $\mathbf{S}_k^\mu$ of locations to obtain measurements $\{\mathbf{y}_k^{(j)}\}_{j \in \mathbf{S}_k^\mu}$ can be computed using (2). Notice that according to (4a), for each selected sampling location j, the system obtains the measurement $\mathbf{x}_k^{(j)} + \mathbf{w}_k^{(j)}$ with probability p or it fails to obtain any measurement with probability $1-p$. We assume that the event of obtaining no measurement is independent of $\mathbf{v}_k^{(j)}, \mathbf{u}_k^{(j)}, \mathbf{w}_k^{(j)}$. □

III. Problem Formulation

A. Notation and Definitions

We begin with describing some of key notation and definitions used throughout the paper.

Notation:

- M – the total number of sensors available for collecting measurements.
- N – the length of the time horizon over which the stochastic process is monitored.
- $\{s_k\}_{k=1}^N$ – positive constants that describe the maximum number s_k of sensors to be selected at each time index k in $\{1, \cdots, N\}$.
- $\mathbf{x}_{1:N}$ – given random variables $\{\mathbf{x}_k\}_{k=1}^N$, $\mathbf{x}_{1:N}$ is shorthand notation for $\mathbf{x}_{1:N} = (\mathbf{x}_1, \cdots, \mathbf{x}_N)$.
- $\mathbf{y}_{\mathbf{S}_{1:N}}$ – given random variables $\{\mathbf{y}_k^{(j)}\}_{j=1}^M$ and a subset $\mathbf{S}_k$ of $\{1, \cdots, M\}$ for k in $\{1, \cdots, N\}$, $\mathbf{y}_{\mathbf{S}_{1:N}}$ is shorthand notation for $\mathbf{y}_{\mathbf{S}_{1:N}} = (\mathbf{y}_{\mathbf{S}_1}, \cdots, \mathbf{y}_{\mathbf{S}_N})$ with $\mathbf{y}_{\mathbf{S}_k} = (\mathbf{S}_k, \{\mathbf{y}_k^{(j)}\}_{j \in \mathbf{S}_k})$.
- $\mathcal{S}_k$, $\mathcal{Y}_k$ – given a positive constant s_k, $\mathcal{S}_k$ represents a collection of subsets $\mathbf{S}_k$ of $\{1, \cdots, M\}$ for which $|\mathbf{S}_k| \leq s_k$ holds, and $\mathcal{Y}_k$ represents a collection of measurements $\mathbf{y}_{\mathbf{S}_k}$ for every $\mathbf{S}_k \in \mathcal{S}_k$.
- $\overline{\mathcal{S}}_k$, $\overline{\mathcal{Y}}_k$ – given a positive constant s_k, $\overline{\mathcal{S}}_k$ represents a collection of subsets $\mathbf{S}_k$ of $\{1, \cdots, M\}$ for which $|\mathbf{S}_k| = s_k$ holds, and $\overline{\mathcal{Y}}_k$ represents a collection of measurements $\mathbf{y}_{\mathbf{S}_k}$ for every $\mathbf{S}_k \in \overline{\mathcal{S}}_k$.

In what follows, we define admissible adaptive policies and information states.

***Definition** III.1 (Admissible Adaptive Policy):* An adaptive policy $\mu = (\mu_1, \cdots, \mu_N)$ is called *admissible* if each μ_k is defined by[3]

$$\mu_k : \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1} \times \mathbb{Z}^k \to \overline{\mathcal{S}}_k \tag{5}$$

We denote the set of all admissible policies by $\mathcal{U}$.

***Remark** III.2:* We call a policy $\mu = (\mu_1, \cdots, \mu_N)$ *adaptive* as each μ_k depends on measurements $\mathbf{y}_{\mathbf{S}_{1:k-1}} \in \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1}$ and auxiliary data $\mathbf{z}_{1:k} \in \mathbb{Z}^k$ in assigning a set $\mathbf{S}_k \in \overline{\mathcal{S}}_k$ as in (5). We refer to a policy μ *non-adaptive* if it only depends on sensor selection sets $\mathbf{S}_{1:k-1} \in \mathcal{S}_1 \times \cdots \times \mathcal{S}_{k-1}$.

***Definition** III.3 (Admissible Information State):* We refer to $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k}) \in \mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1} \times \mathbb{Z}^k$ as an *admissible information state* at time index k if there is an admissible policy $\mu = (\mu_1, \cdots, \mu_N)$ for which $\mathbf{S}_j = \mu_j(\mathbf{y}_{\mathbf{S}_{1:j-1}}, \mathbf{z}_{1:j})$ holds for all j in $\{1, \cdots, k-1\}$.

***Remark** III.4 (Sensor Selection Subset):* Given an admissible information state $(\mathbf{y}_{\mathbf{S}_{1:N}}, \mathbf{z}_{1:N})$ with its associated policy $\mu = (\mu_1, \cdots, \mu_N) \in \mathcal{U}$, we represent each sensor selection set $\mathbf{S}_k$ by $\mathbf{S}_k = \{\mathbf{j}_1, \cdots, \mathbf{j}_{s_k}\}$. The i-th element $\mathbf{j}_i$ can be considered as a random variable defined by $\mathbf{j}_i = \mu_k^i(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$, where μ_k^i is the i-th component of μ_k.[4] We refer to $\breve{\mathbf{S}}_k$ as a *subset* of $\mathbf{S}_k$ if it holds that $\breve{\mathbf{S}}_k = \{\mathbf{j}_{i_1}, \cdots, \mathbf{j}_{i_{s'_k}}\}$ with $\{i_1, \cdots, i_{s'_k}\} \subset \{1, \cdots, s_k\}$ and $s'_k < s_k$.

***Remark** III.5:* Note that given an admissible information state $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$ and an admissible policy $\mu = (\mu_1, \cdots, \mu_N)$, we can derive

$$\mathbf{S}_j^\mu = \mu_j(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{y}_{\mathbf{S}^\mu_{k:j-1}}, \mathbf{z}_{1:j}) \tag{6}$$

for all j in $\{k, \cdots, N\}$. We adopt the shorthand notation $\mathbf{S}^\mu_{k:N} = \mu_{k:N}(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$ to denote the sensor selection sets $\{\mathbf{S}_j^\mu\}_{j=k}^N$ determined by μ according to (6) using $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$.

B. Problem Description

We formally state the main problem of designing an adaptive policy for sensor selection.

***Problem** III.6:* Find an adaptive policy $\mu = (\mu_1, \cdots, \mu_N)$ in $\mathcal{U}$ that maximizes

$$\mathcal{J}(\mu) = \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\mathbf{S}^\mu_{1:N}}, \mathbf{z}_{1:N}) \tag{7}$$

subject to

$$\mathbf{S}_k^\mu = \mu_k(\mathbf{y}_{\mathbf{S}^\mu_{1:k-1}}, \mathbf{z}_{1:k}) \tag{8a}$$

$$|\mathbf{S}_k^\mu| \leq s_k \tag{8b}$$

[3]Note that without loss of optimality, we will be restricting our attention to adaptive policies that assign sensor selection sets with the maximum cardinality. This is mainly because (1) only increases with more sensor measurements.

[4]Here the order in labeling the elements of μ_k can be arbitrary and our results do not depend on the ordering.

Note that using the chain rule of MI, we can see that (1) and (7) are different by a constant which does not depend on the choice of a policy.

We maintain the following assumption throughout the paper.

***Assumption** III.7:* We assume that the joint probability density function

$$p(\mathbf{x}_1, \cdots, \mathbf{x}_N, \mathbf{y}_1^{(1)}, \cdots, \mathbf{y}_N^{(M)}, \mathbf{z}_1, \cdots, \mathbf{z}_N) \tag{9}$$

is well-defined, and given $\mathbf{x}_{1:N}$, the random variables $\mathbf{y}_k^{(j)}$ and $\mathbf{z}_k$ for all j in $\{1, \cdots, M\}$ and k in $\{1, \cdots, N\}$ are mutually independent.

In what follows, we state two conditions on adaptive policies, which will be used to derive the main results of this paper. To this end, let us re-write the performance index (7) as follows: For each k in $\{1, \cdots, N+1\}$,

$$\begin{aligned}\mathcal{J}(\mu) = &\mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\mathbf{S}^\mu_{1:k-1}}, \mathbf{z}_{1:k}) \\ &+ \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\mathbf{S}^\mu_{k:N}}, \mathbf{z}_{k+1:N} \,|\, \mathbf{y}_{\mathbf{S}^\mu_{1:k-1}}, \mathbf{z}_{1:k})\end{aligned} \tag{10}$$

with $\mathbf{z}_{N+1} = 0$. The second term in (10) can be viewed as the cost-to-go of (7) from time index k given the information state $(\mathbf{y}_{\mathbf{S}^\mu_{1:k-1}}, \mathbf{z}_{1:k})$. In terms of this cost-to-go, we define diminishing property and (necessary) optimality conditions as follows.

***Definition** III.8 (Diminishing Property Condition):* We say that an admissible policy $\mu = (\mu_1, \cdots, \mu_N) \in \mathcal{U}$ satisfies the *diminishing property condition* if the following inequality holds for any admissible information state $(\mathbf{y}_{\mathbf{S}_{1:N-1}}, \mathbf{z}_{1:N})$: For each k in $\{1, \cdots, N\}$,

$$\begin{aligned}&\mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\mathbf{S}^\mu_{k:N}}, \mathbf{z}_{k+1:N} \,|\, \mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\mathbf{S}_{k-1}}, \mathbf{z}_{1:k}) \\ &\leq \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\breve{\mathbf{S}}^\mu_{k:N}}, \mathbf{z}_{k+1:N} \,|\, \mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\breve{\mathbf{S}}_{k-1}}, \mathbf{z}_{1:k})\end{aligned} \tag{11}$$

where $\mathbf{S}^\mu_{k:N} = \mu_{k:N}(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\mathbf{S}_{k-1}}, \mathbf{z}_{1:k})$ and $\breve{\mathbf{S}}^\mu_{k:N} = \mu_{k:N}(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\breve{\mathbf{S}}_{k-1}}, \mathbf{z}_{1:k})$ for an arbitrary subset $\breve{\mathbf{S}}_{k-1}$ of $\mathbf{S}_{k-1}$.

***Definition** III.9 (Optimality Condition):* We say that an admissible policy $\mu = (\mu_1, \cdots, \mu_N) \in \mathcal{U}$ satisfies the *optimality condition* if the following inequality holds for any admissible information states $(\mathbf{y}_{\mathbf{S}_{1:N-1}}, \mathbf{z}_{1:N})$: For each k in $\{1, \cdots, N\}$,

$$\begin{aligned}&\mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\breve{\mathbf{S}}^\mu_{k:N}}, \mathbf{z}_{k+1:N} \,|\, \mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k}) \\ &\leq \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_{\mathbf{S}^\mu_{k:N}}, \mathbf{z}_{k+1:N} \,|\, \mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})\end{aligned} \tag{12}$$

where $\mathbf{S}^\mu_{k:N} = \mu_{k:N}(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\mathbf{S}_{k-1}}, \mathbf{z}_{1:k})$ and $\breve{\mathbf{S}}^\mu_{k:N} = \mu_{k:N}(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\breve{\mathbf{S}}_{k-1}}, \mathbf{z}_{1:k})$ for an arbitrary subset $\breve{\mathbf{S}}_{k-1}$ of $\mathbf{S}_{k-1}$.

An intuition behind the diminishing property condition is that the cost-to-go in (10) attains a higher value when it is assessed conditioned on the smaller information state $(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\breve{\mathbf{S}}_{k-1}}, \mathbf{z}_{1:k})$. The notion of the diminishing property for adaptive policies is closely related with submodularity of set functions used in combinatorial optimization literature [9]. We will briefly discuss their relationship in Section V.

Algorithm 1: Greedy Adaptive Policy μ_k^g

input : Information state $(y_{S_{1:k-1}}, z_{1:k})$
output: Sensor Selection Set $S_k^g = \mu_k^g\left(y_{S_{1:k-1}}, z_{1:k}\right)$

1 **begin**
2 $\quad S_k^g \leftarrow \varnothing$
3 $\quad$ **while** $|S_k^g| < s_k$ **do**
4 $\quad\quad j^g \leftarrow \underset{j \in \{1, \cdots, M\} \setminus S_k^g}{\arg\max} \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_k^{(j)} \,|\, \mathbf{z}_{1:k} = z_{1:k},$
5 $\quad\quad\quad \mathbf{y}_{\mathbf{S}_{1:k-1}} = y_{S_{1:k-1}}, \mathbf{y}_{S_k^g})$
6 $\quad\quad S_k^g \leftarrow S_k^g \cup \{j^g\}$

The optimality condition has the following implication: The policy $\mu = (\mu_1, \cdots, \mu_N)$ assigns sensor selection sets $\mathbf{S}_{k:N}^{\mu}$ with the information state $(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\mathbf{S}_{k-1}}, \mathbf{z}_{1:k})$ that attain a higher cost-to-go (12) than do other sensor selection sets $\breve{\mathbf{S}}_{k:N}^{\mu}$ determined by the same policy but with the *smaller* information state $(\mathbf{y}_{\mathbf{S}_{1:k-2}}, \mathbf{y}_{\breve{\mathbf{S}}_{k-1}}, \mathbf{z}_{1:k})$. Associated with Definition III.9, notice that if a policy $\mu \in \mathcal{U}$ does not satisfy (12), then μ cannot be the optimal policy.

C. Greedy Adaptive Policies

Finding the optimal adaptive policy is NP-hard even for a special case of our problem [1]. To find a computationally tractable solution for Problem III.6, we adopt the greedy heuristics to our problem setting and define greedy adaptive policies as follows.

Definition III.10 (Greedy Adaptive Policy): An adaptive policy $\mu^g = (\mu_1^g, \cdots, \mu_N^g) \in \mathcal{U}$ is called *greedy* if for every element $(y_{S_{1:k-1}}, z_{1:k})$ in $\mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1} \times \mathbb{Z}^k$, each μ_k^g assigns a sensor selection set S_k^g according to Algorithm 1.

One key feature of Algorithm 1 is that it does not require any look-ahead in computing the sensor selection set S_k^g: Algorithm 1 consists of s_k rounds of sensor selection (Line 4) in which each round only requires to compute MI between $\mathbf{x}_{1:N}$ and $\mathbf{y}_k^{(j)}$ given a realization $(y_{S_{1:k-1}}, z_{1:k})$ and the set S_k^g representing sensors selected in prior rounds.

IV. Related Work

The problem of designing sensor selection policies for monitoring stochastic processes is closely related with sensor scheduling/management problems investigated across multiple disciplines.

A number of efforts have been made to develop optimization frameworks to find optimal sensor selection schemes. A convex relaxation method for a sensor scheduling problem was developed to estimate a vector parameter [10] and a stochastic signal [11]. The work of [12] investigated a stochastic sensor scheduling problem using a convex optimization framework. More recently [13] presented a relaxed convex formulation for a sparse sensor selection problem. For linear dynamical systems, tree representation-based methods and pruning-based computational methods to find optimal solutions are developed in [14], [15]. In addition, the POMDP framework was adopted in a sensor scheduling problem [16] for optimal scheduler design.

However, finding the optimal solution is in general NP-complete [17], and optimization-based approaches require substantial computational resources. The paper of [9] investigated the notion of submodularity, which was originally introduced in combinatorial optimization literature [2], [18], in resource allocation problems, and showed that greedy heuristics provide computationally feasible and sub-optimal solutions in submodular problems. As applications of submodularity methods to sensor scheduling problems with information theoretic performance indices, the work of [4] verified that MI is submodular and the greedy heuristics ensure the performance of greedy sensor scheduling schemes. On the other hand, log estimation error ellipsoid is shown to be submodular in [19]; while the authors of [20] investigated submodularity of various performance indices used in sensor scheduling problems. For further reading, the reader is referred to survey papers [21], [22] and references therein.

While the above-mentioned work mainly focuses on sensor selection problems in non-adaptive settings, the work of [23] investigated a stronger notion of submodularity, namely adaptive submodularity, to investigate the performance of adaptive greedy heuristics in learning and optimization problem domains. An application of adaptive sumbodularity in batch mode active learning problems is found in [24]. Analytical tools for adaptive sensor selection problems are further developed in [25], [26] to establish the performance gap between non-adaptive and adaptive policies [25], or to explore sensor scheduling problems in continuous space domains [26].

V. Performance Analysis of Greedy Adaptive Policies

In this section, we analyze the performance of the greedy adaptive policy and establish a connection to results from submodular optimization literature.

A. Performance Analysis

We first proceed with performance analysis on greedy adaptive policies.

Theorem V.1: Let $\mathcal{U}^D$ be the set of admissible adaptive policies satisfying Definitions III.8 and III.9. For a greedy adaptive policy μ^g, it holds that

$$\frac{1}{2} \max_{\mu \in \mathcal{U}^D} \mathcal{J}(\mu) \leq \mathcal{J}(\mu^g)$$

In general, the optimal adaptive policy in $\mathcal{U}$ does not belong to $\mathcal{U}^D$, and thus the performance bound on the greedy adaptive policy cannot be directly related with the performance of the optimal policy. However, depending on applications and problem domains, important classes of adaptive policies are associated with $\mathcal{U}^D$.[5]

Let us define a set $\mathcal{U}^z$ of admissible adaptive policies $\mu^z = (\mu_1^z, \cdots, \mu_N^z)$ that assign sensor selection sets $\mathbf{S}_{1:N}^{\mu^z}$

[5] Due to space restrictions, we provide a few examples in this paper and leave in-depth analysis for future work.

as follows: For each k in $\{1, \cdots, N\}$, given $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$ in $\mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1} \times \mathbb{Z}^k$,

$$\mathbf{S}_k^{\mu^z} = \mu_k^z(\mathbf{S}_{1:k-1}, \mathbf{z}_{1:k}) \tag{13}$$

where $\mathbf{S}_k^{\mu^z}$ belongs to $\overline{\mathcal{S}}_k$. Note that the policy $\mu^z = (\mu_1^z, \cdots, \mu_N^z)$ depends on the choice of sensor selection sets $\mathbf{S}_{1:k-1}$ and auxiliary data $\mathbf{z}_{1:k}$.

The following proposition states that the optimal policy in $\mathcal{U}^z$ belongs to $\mathcal{U}^D$.

Proposition V.2: Consider a class $\mathcal{U}^z$ of admissible policies $\mu^z = (\mu_1^z, \cdots, \mu_N^z)$ that assign sensor selection sets $\mathbf{S}_{1:N}^{\mu^z}$ according to (13). The optimal policy in $\mathcal{U}^z$ belongs to $\mathcal{U}^D$.

For the models given in Example II.1, let us define a set $\mathcal{U}^p$ of admissible adaptive policies $\mu^p = (\mu_1^p, \cdots, \mu_N^p)$ that assign sensor selection sets $\mathbf{S}_{1:N}^{\mu^p}$ as follows: For each k in $\{1, \cdots, N\}$, given $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$ in $\mathcal{Y}_1 \times \cdots \times \mathcal{Y}_{k-1} \times \mathbb{Z}^k$,

$$\mathbf{S}_k^{\mu^p} = \mu_k^p(\mathbf{S}_{1:k-1}, \mathbf{\Phi}_{1:k-1}, \mathbf{z}_{1:k}) \tag{14}$$

where $\mathbf{S}_k^{\mu^p}$ belongs to $\overline{\mathcal{S}}_k$ and $\mathbf{\Phi}_l \overset{\text{def}}{=} \{j \in \mathbf{S}_l \,|\, \mathbf{y}_l^{(j)} = \varnothing\}$ is a set of selected sensors from which the system collects $\varnothing$, i.e., the system selects sensor $j \in \mathbf{\Phi}_l$, but it fails to collect any measurement.

The following proposition states that the optimal policy in $\mathcal{U}^p$ belongs to $\mathcal{U}^D$.

Proposition V.3: Consider the models described in (3) and (4) given in Example II.1 and a class $\mathcal{U}^p$ of admissible policies $\mu^p = (\mu_1^p, \cdots, \mu_N^p)$ that assign sensor selection sets $\mathbf{S}_{1:N}^{\mu^p}$ according to (14). The optimal policy in $\mathcal{U}^p$ belongs to $\mathcal{U}^D$.

B. Connection to Submodularity

In combinatorial optimization literature, the notion of submodularity has been served as an important condition under which greedy heuristics provide sub-optimal solutions. However, MI is in general not submodular in adaptive sensor selection problems, and assessing the performance of greedy heuristics-based approaches is challenging. In comparison with results in the literature, our main findings have the following advantages.

1) Our main result – Theorem V.1 – applies to the adaptive sensor selection problem without submodularity requirements and along with Propositions V.2 and V.3, it provides a non-trivial lower bound on the performance of greedy policies.
2) Notably if the performance index (7) satisfies the adaptive submodularity condition [23], then the optimal adaptive policy in $\mathcal{U}$ belongs to $\mathcal{U}^D$, and based on Theorem V.1, we can derive the performance bound on greedy adaptive policies with respect to the performance of the optimal adaptive policy.

To explain the second argument, we adopt the adaptive submodularity condition from [23], and establish the performance bound on greedy adaptive policies, which extends

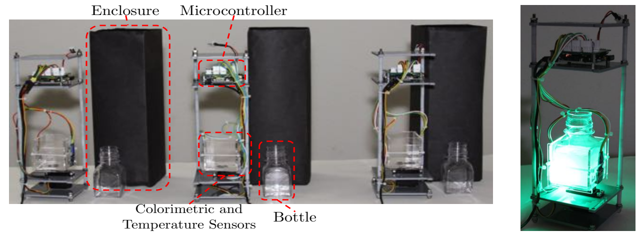

(a) (b)

Fig. 1. (a) The test bed consisting of 3 sensor nodes. (b) Absorbance measurement using the colorimetric sensor.

results of [4] to adaptive sensor selection problems.[6] For simplicity, we consider the case where $s_k = 1$ for all k in $\{1, \cdots, N\}$.

Definition V.4: The performance index (7) is said to be *adaptive submodular* if the following condition holds for all j in $\{1, \cdots, M\}$ and k in $\{1, \cdots, N\}$: Given any admissible information state $(\mathbf{y}_{\mathbf{S}_{1:N-1}}, \mathbf{z}_{1:N})$, for each realization $(y_{S_{1:k-1}}, z_{1:k})$ of $(\mathbf{y}_{\mathbf{S}_{1:k-1}}, \mathbf{z}_{1:k})$, it holds that

$$\begin{aligned} &\mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_k^{(j)}, \mathbf{z}_{k+1} \,|\, \mathbf{y}_{\mathbf{S}_{1:k-1}} = y_{S_{1:k-1}}, \mathbf{z}_{1:k} = z_{1:k}) \\ &\leq \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_k^{(j)}, \mathbf{z}_{k+1} \,|\, \mathbf{y}_{\breve{\mathbf{S}}_{1:k-1}} = y_{\breve{S}_{1:k-1}}, \mathbf{z}_{1:k} = z_{1:k}) \end{aligned} \tag{15}$$

where, in (15), $\breve{\mathbf{S}}_l$ is an arbitrary subset of $\mathbf{S}_l$ and $\breve{S}_l$ is a subset of the realization S_l of $\mathbf{S}_l$ for all l in $\{1, \cdots, k-1\}$.

According to Definition V.4, it can be verified that the optimal adaptive policy in $\mathcal{U}$ satisfies the diminishing property condition. Based on this observation, we can state the following corollary.

Corollary V.5: Suppose that the performance index (7) is adaptive submodular. For a greedy adaptive policy μ^g, it holds that

$$\frac{1}{2} \max_{\mu \in \mathcal{U}} \mathcal{J}(\mu) \leq \mathcal{J}(\mu^g)$$

The proof directly follows from Definitions III.8 and V.4 and Theorem V.1.

Note that the adaptive submodularity is a condition for the performance index whereas the diminishing property and optimality conditions are defined for policies. Corollary V.5 implies that if the performance index is adaptive submodular, then the optimal adaptive policy belongs to $\mathcal{U}^D$.

VI. Experiments Results

In this section, we demonstrate how the adaptive sensor selection scheme can be applied to the motivating application. For this purpose, we use our test bed consisting of 3 wireless sensor nodes and perform experiments on monitoring the population growth of yeast. Each node in the test bed is equipped with a WiFi-enabled microcontroller and colorimetry sensor (see Fig. 1).

Yeast Population Growth: Yeast is a type of a fungus consisting of single cells that reproduce via a process called budding. The population growth of yeast can be measured by evaluating how much a medium containing yeast absorbs

[6]Note that the problem considered in [23] is different from Problem III.6, and results from [23] cannot be used to conclude our main results.

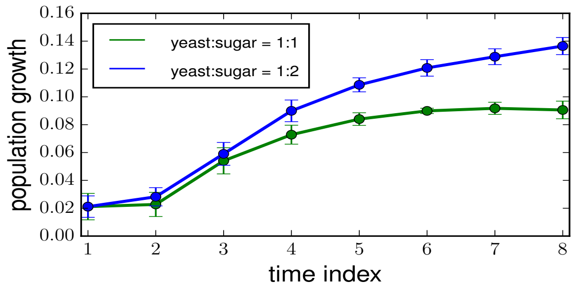

Fig. 2. Population growth data of yeast at two different ratios of yeast and sugar. The data points are measured at $30 - \mathrm{min}$ intervals.

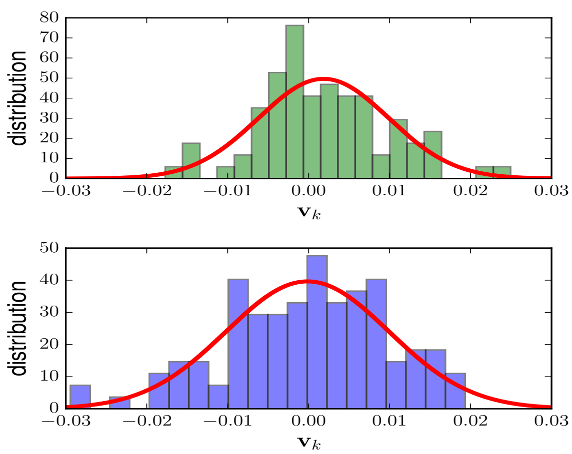

Fig. 3. The histograms depicting the model error computed using the experiment data and the graphs of the probability distributions of $\mathbf{v}_k^{(j)}$ with the parameters: (upper figure – yeast : sugar $= 1:1$) $\sigma_{\mathbf{v}}^{(j)} = 0.0089$, (lower figure – yeast : sugar $= 1:2$) $\sigma_{\mathbf{v}}^{(j)} = 0.0093$.

light at a certain wavelength (Beer's law). For this reason, we will use absorbance to denote the yeast population and adopt the model (3) to describe the population growth, where the noise term $\mathbf{v}_k^{(j)}$ is assume to be zero-mean Gaussian and to have a constant variance $\sigma_{\mathbf{v}}^{(j)}$ over entire time horizon.

In experiments, we use active dry yeast dissolved in water with its food source (granulated sugar) at two different ratios – $1:1$ and $1:2$ (see Fig. 2 for population growth data). For each of these cases, the growth rate $r^{(j)}$, carrying capacity $K^{(j)}$, and the variance $\sigma_{\mathbf{v}}^{(j)}$ of $\mathbf{v}_k^{(j)}$ can be computed from experiment data using the maximum likelihood estimation (MLE). Fig. 3 illustrates the results on estimating $\sigma_{\mathbf{v}}^{(j)}$.

Yeast Samples and Population Measurements: During each experiment, yeast samples are collected at $30 - \mathrm{min}$ intervals for 8 times ($N = 8$). Collected yeast samples are processed using lab-based equipment to obtain population measurements. We use the model (4a) to describe how the measurement depends on the yeast population. For simplicity, we assume that the probability of failing to collect yeast samples is zero, i.e., $p = 1$ for (4a), and that there is no disturbance, i.e., $\mathbf{w}_k^{(j)} = 0$.

Sensor Data: The colorimetric sensor is implemented

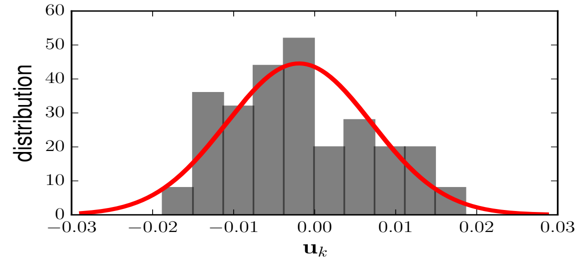

Fig. 4. The histogram depicting the colorimetric sensing error and the graph of the probability distribution of $\mathbf{u}_k^{(i)}$ with the parameter $\sigma_{\mathbf{u}}^{(i)} = 0.009$.

TABLE I

PARAMETERS SELECTED FOR EXPERIMENTS

parameters	values	parameters	values	parameters	values
$r^{(1)}$	1.4	$r^{(2)}$	1.2	$r^{(3)}$	1.2
$K^{(1)}$	0.085	$K^{(2)}$	0.130	$K^{(3)}$	0.130
$\sigma_{\mathbf{v}}^{(1)}$	0.0089	$\sigma_{\mathbf{v}}^{(2)}$	0.0093	$\sigma_{\mathbf{v}}^{(3)}$	0.0093
$\sigma_{\mathbf{u}}^{(1)}$	0.009	$\sigma_{\mathbf{u}}^{(2)}$	0.009	$\sigma_{\mathbf{u}}^{(3)}$	0.009
$\sigma_{\mathbf{w}}^{(1)}$	0	$\sigma_{\mathbf{w}}^{(2)}$	0	$\sigma_{\mathbf{w}}^{(3)}$	0

at every node and is used to collect noisy measurements of yeast population. The sensor consists of a green LED (527 nm) and a photodiode array (TCS3472, AMS): The LED is used as a light source and the photodiode array measures the amount of light arrives from the source (see Fig. 1(b)). From the output of the photodiode array, we compute the absorbance using Beer's law. We use the model (4b) to describe how the sensor data depends on the yeast population, where the disturbance $\mathbf{u}_k^{(j)}$ is assumed to be zero-mean Gaussian and to have a constant variance $\sigma_{\mathbf{u}}^{(j)}$ over entire time horizon. From experiment data, we compute the variance $\sigma_{\mathbf{u}}^{(j)}$ of $\mathbf{u}_k^{(j)}$ using the MLE. Fig. 4 illustrates the results on estimating $\sigma_{\mathbf{u}}^{(j)}$.

Experiment Protocol: To monitor the population growth of yeast growing in 3 different bottles ($M = 3$), we use 3 sensor nodes each of which is deployed with one of the yeast bottles. The yeast population is monitored every 30 mins for 8 times ($N = 8$): At each time index $k \in \{1, \cdots, 8\}$, the colorimetry sensors collect noisy measurements $\mathbf{z}_k$ of yeast population $\mathbf{x}_k$. Based on the greedy policy μ, at each time index $k \in \{1, \cdots, 8\}$, we collect a yeast sample from one bottle, say bottle j, (out of 3 bottles) to obtain an accurate population measurement $\mathbf{y}_k^{(j)}$ of the yeast in bottle j using lab-based equipment. The experiments are repeated for 5 times. Table I summarizes the parameters for the models (3) and (4) used in the experiments.

Experiment Results and Performance Analysis: To assess the performance of the greedy adaptive policy μ, at each time index k in $\{1, \cdots, 8\}$, we compute the conditional entropy $h(\mathbf{x}_{1:8} \,|\, \mathbf{y}_{\mathbf{S}^{\mu}_{1:k-1}} = y_{S^{\mu}_{1:k}}, \mathbf{z}_{1:k} = z_{1:k})$ given the population measurements $(y_{S^{\mu}_{1:k}}, z_{1:k})$, where $\mathbf{x}_k = (\mathbf{x}_k^{(1)}, \mathbf{x}_k^{(2)}, \mathbf{x}_k^{(3)})$ and each set S_k^{μ} is selected according to Algorithm 1.[7] The

[7]Note that given the measurement $(y_{S^{\mu}_{1:8}}, z_{1:8})$, the performance index (7) can be represented as $\mathcal{J}(\mu) = h(\mathbf{x}_{1:8}) - h(\mathbf{x}_{1:8} \,|\, \mathbf{y}_{\mathbf{S}^{\mu}_{1:8}} = y_{S^{\mu}_{1:8}}, \mathbf{z}_{1:8} = z_{1:8})$.

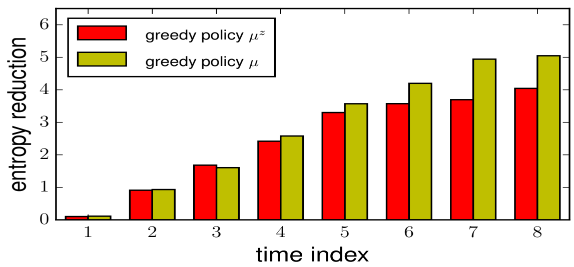

Fig. 5. The graph illustrates the entropy reduction for each k in $\{1, \cdots, 8\}$ with policy μ and policy μ^z.

conditional entropy represents the uncertainty in the yeast population $\mathbf{x}_{1:8}$ over the entire experiment period given the measurement $(y_{S^{\mu}_{1:k}}, z_{1:k})$. Fig. 5 depicts the reduction in the conditional entropy with respect to $h(\mathbf{x}_{1:8})$, averaged over the 5-experiment trials.

In addition to illustrating the performance of the policy μ determined by Algorithm 1, we compute another greedy policy μ^z from the set $\mathcal{U}^z$ (see (13)) to see how the performance of greedy-based heuristics changes based on the amount of information that the heuristics depend on. The greedy policy μ^z assigns each set $S_k^{\mu^z}$ according to Algorithm 1 where Line 4 is substituted with

$$j^g \leftarrow \underset{j \in \{1, \cdots, M\} \setminus S_k^{\mu^z}}{\arg\max} \mathcal{I}(\mathbf{x}_{1:N}; \mathbf{y}_k^{(j)} \,|\, \mathbf{y}_{S_{1:k-1}}, \mathbf{y}_{S_k^{\mu^z}}, \mathbf{z}_{1:k} = z_{1:k})$$

To assess the performance of μ^z, we compute the conditional entropy

$$h(\mathbf{x}_{1:8} \,|\, \mathbf{y}_{S_{1:k}^{\mu^z}}, \mathbf{z}_{1:k} = z_{1:k})$$

given the sampling location sets and sensor data $(S_{1:k}^{\mu^z}, z_{1:k})$. Fig. 5 depicts the reduction in the conditional entropy with respect to $h(\mathbf{x}_{1:8})$, averaged over the 5-experiment trials. Notice that the greedy policy μ, which depends on both the measurements and sensor data, outperforms the greedy policy μ^z which only depends on the sensor data.

VII. Conclusions

We have investigated designing adaptive policies in an adaptive sensor selection problem. The main contributions of this paper are to define greedy adaptive policies and establish the performance bound on greedy policies. For this purpose, we have defined a particular class $\mathcal{U}^D$ of policies satisfying the diminishing property and optimality conditions, and showed that the performance of greedy policies is lower bounded by $\frac{1}{2}$ of that of the best policy within the class $\mathcal{U}^D$. As future work, we will identify types of adaptive policies belonging to $\mathcal{U}^D$ with respect to various stochastic models.

References

[1] V. Tzoumas, N. A. Atanasov, A. Jadbabaie, and G. J. Pappas, "Scheduling nonlinear sensors for stochastic process estimation," *CoRR*, vol. abs/1609.08536, 2016.

[2] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, *An analysis of approximations for maximizing submodular set functions—II*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 73–87.

[3] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," in *Proceedings of the 28th International Conference on International Conference on Machine Learning*, ser. ICML'11. USA: Omnipress, 2011, pp. 1057–1064.

[4] J. L. Williams, J. W. F. III, and A. S. Willsky, "Performance guarantees for information theoretic active inference," *Journal of Machine Learning Research*, vol. 2, pp. 620–627, 2007.

[5] "Underworlds," http://underworlds.mit.edu.

[6] C. L. Kaiser, D. R. Yoerger, J. C. Kinsey, S. Kelley, A. Billings, J. Fujii, S. Suman, M. Jakuba, Z. Berkowitz, C. R. German, and A. D. Bowen, "The design and 200 day per year operation of the autonomous underwater vehicle Sentry," in *2016 IEEE/OES Autonomous Underwater Vehicles (AUV)*, Nov 2016, pp. 251–260.

[7] D. Diamond, F. Collins, J. Cleary, C. Zuliani, and C. Fay, *Distributed Environmental Monitoring*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 321–363.

[8] F. Brauer and C. Castillo-Chavez, *Discrete Population Models*. New York, NY: Springer New York, 2012, pp. 49–90.

[9] A. Krause and C. Guestrin, "Submodularity and its applications in optimized information gathering," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 4, pp. 32:1–32:20, July 2011.

[10] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Transactions on Signal Processing*, vol. 57, no. 2, pp. 451–462, Feb 2009.

[11] Y. Mo, R. Ambrosino, and B. Sinopoli, "Sensor selection strategies for state estimation in energy constrained wireless sensor networks," *Automatica*, vol. 47, no. 7, pp. 1330 – 1338, 2011.

[12] C. Li and N. Elia, "Stochastic sensor scheduling via distributed convex optimization," *Automatica*, vol. 58, pp. 173 – 182, 2015.

[13] V. Roy, A. Simonetto, and G. Leus, "Spatio-temporal sensor management for environmental field estimation," *Signal Processing*, vol. 128, pp. 369 – 381, 2016.

[14] A. Logothetis and A. Isaksson, "On sensor scheduling via information theoretic criteria," in *Proceedings of the 1999 American Control Conference (Cat. No. 99CH36251)*, vol. 4, 1999, pp. 2402–2406 vol.4.

[15] M. P. Vitus, W. Zhang, A. Abate, J. Hu, and C. J. Tomlin, "On efficient sensor scheduling for linear dynamical systems," *Automatica*, vol. 48, no. 10, pp. 2482 – 2493, 2012.

[16] V. Krishnamurthy, "POMDP sensor scheduling with adaptive sampling," in *17th International Conference on Information Fusion (FUSION)*, July 2014, pp. 1–7.

[17] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, June 2008.

[18] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—i," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.

[19] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *49th IEEE Conference on Decision and Control (CDC)*, Dec 2010, pp. 2572–2577.

[20] S. T. Jawaid and S. L. Smith, "Submodularity and greedy algorithms in sensor scheduling for linear dynamical systems," *Automatica*, vol. 61, pp. 282 – 288, 2015.

[21] A. O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3064–3075, Dec 2011.

[22] H. Rowaihy, S. Eswaran, M. Johnson, D. Verma, A. Bar-noy, and T. Brown, "A survey of sensor selection schemes in wireless sensor networks," in *In SPIE Defense and Security Symposium Conference on Unattended Ground, Sea, and Air Sensor Technologies and Applications IX*, 2007.

[23] D. Golovin and A. Krause, "Adaptive submodularity: Theory and applications in active learning and stochastic optimization," *J. Artif. Int. Res.*, vol. 42, no. 1, pp. 427–486, Sept. 2011.

[24] Y. Chen and A. Krause, "Near-optimal batch mode active learning and adaptive submodular optimization," in *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, S. Dasgupta and D. Mcallester, Eds. JMLR Workshop and Conference Proceedings, 2013, pp. 160–168.

[25] A. Asadpour and H. Nazerzadeh, "Maximizing stochastic monotone submodular functions," *Management Science*, vol. 62, no. 8, pp. 2374–2391, 2016.

[26] A. A. Bian, B. Mirzasoleiman, J. M. Buhmann, and A. Krause, "Guaranteed non-convex optimization: Submodular maximization over continuous domains," *CoRR*, vol. abs/1606.05615, 2016.