

Online Decision Making for Stream-based Robotic Sampling via Submodular Optimization

Wenhao Luo, *Student Member, IEEE*, Changjoo Nam, *Member, IEEE*, Katia Sycara, *Fellow, IEEE*

Abstract—We consider the problem of online robotic sampling in environmental monitoring tasks where the goal is to collect k best samples from n sequentially occurring measurements. In contrast to many existing works that seek to maximize the utility of the selected samples online, we aim to find the cardinality constrained subset of streaming measurements under irrevocable sampling decisions so that the prediction over untested measurements is most accurate. Using the information theoretic criterion, we present an online submodular algorithm for stream-based sample selection with a provable performance bound. We demonstrate the effectiveness of our algorithm via simulations of information gathering from indoor static sensors.

I. INTRODUCTION

Robotic sampling has received extensive attention in field applications such as ecosystem monitoring [1], crop phenotyping [2], environmental surveillance [3], and patrolling [4]. Samples are needed so as to model and infer the hidden distribution of features of interest. In many cases the sampling cost (e.g., time, on-board sample storage capacity, or energy consumption) scales up quickly, therefore there is a need for choosing samples selectively to collect the most informative sample subset to reduce the prediction uncertainty of the distribution modelling process, known as the *Bayesian experiment design*. There has been prior work addressing the optimal Bayesian experiment design problem, such as [5], [6] where the entire set of available samples are given in advance. However, in many applications, the nature of the samples is not known in advance but samples are presented online in a sequential manner that requires an instantaneous and (often) irrevocable decision of taking the sample or not, known as the *secretary* or *stream-based setting* [7], [8]. To the best of our knowledge, our work presents the first approach to tackle the stream-based Bayesian experiment design problem in robotic sampling.

In this paper, we consider the problem of online decision making for robotic sample selection under the stream-based setting where the robot has a budget limit for sampling (e.g., time limit, battery life, or sample container size). In this problem, a subset of available samples is recursively selected so that the uncertainty from the learned Gaussian process (GP) model from the subset is minimized with a provable bound based on an information theoretic criterion such as mutual information gain. Since the general sample

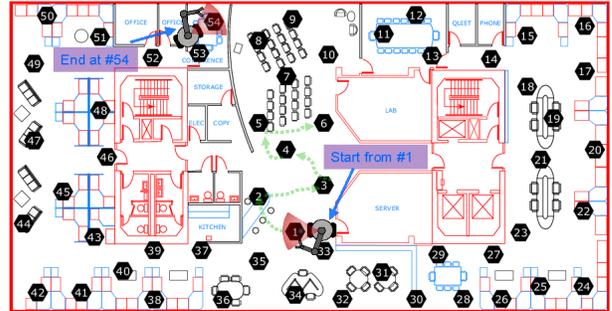


Fig. 1: The robot travels along a fixed path that sequentially visits all 54 deployed stationary sensors in the Intel Berkeley Lab [9] in ascending order of sensor IDs and obtains sensor reading information such as sensor location, humidity, light and voltage to infer the temperature distribution.

selection problem under the cardinality constraint is NP-hard [6], tractable approaches with sub-optimal bound are acceptable.

A testing example scenario is shown in Fig. 1 in which a robot sequentially visits all deployed 54 sensors along a fixed path and obtains the sensor readings including sensor locations, humidity, light and voltage as available features at each corresponding discrete sensor location. We assume that temperature readings are not available during the online sample selection process. Obtaining sensor readings requires additional visual and computing processes consuming battery power (or time in a time-critical data collection), so the number of sensor readings obtained should be constrained. Therefore, the goal is to select the best k sensor readings out of n by deciding whether or not to obtain the sensor data upon visiting each sensor, such that the trained GP model based on the available features could provide good estimation of an objective feature (temperature in our scenario) with minimized uncertainty over the testing sensor data, i.e. the data from the remaining sensors that were not selected. As the value of humidity, light and voltage for those candidate sensors can never be known beforehand, we cannot directly apply the Bayesian experimental design or path planning for sampling algorithms [6], [10]. Therefore, a novel decision making strategy regarding sample selection with the stream-based setting is needed.

Our paper presents two contributions. First, our approach embeds the online sample selection process into a stream-based submodular optimization framework, with a provable performance bound, for better learning the GP model with constrained subset of training data. Second, we provide extensive experimental results using a public dataset [9] from Intel Berkeley Research Lab. Our experiments provide an extensive comparison using different sample selection

This work was supported in part by U.S. Department of Agriculture (DOA) award 2017-6700-726152.

*The authors are with the Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA. Email: {wenhaol, changjon}@andrew.cmu.edu, katia@cs.cmu.edu.

algorithms with different information theoretic criteria and shows the superior results of our approach.

II. RELATED WORK

Using GP models for modelling spatial phenomena has been widely studied for modelling the hidden mapping from available features of training data to the objective features with consideration of uncertainty [11]. Built on the GP modelling, approaches have been proposed to address both the problems of selecting samples with highest utilities (i.e., Bayesian optimization [12]) and learning the underlying distribution of the mapping with as few samples as possible from a given pool of samples known in advance (i.e., Bayesian experiment design [5]).

The problem where (a) the pool of samples are not known in advance, (b) samples are sequentially presented, and (c) the sampling decision is irrevocable is known as *the secretary problem* and has been extensively studied [13], [14]. The classical secretary problem seeks irrevocable online selection of the best candidate (or best k candidates) out of n applicants. The work in [15] considers a variant of the k -choice problem in which the objective is to choose k applicants maximizing the expectation of a *submodular* utility function, which characterizes overlapping skills among the chosen applicants. Das et al. [1] presented the first online algorithm to select k best samples from n candidate measurements in the stream-based setting for robotic monitoring, in which each measurement is evaluated by the utility function specified by the fixed weighted sum of posterior mean and variance functions [12] from an initial pilot survey. Although [1] could select the samples with highest utility in stream-based setting, it does not incorporate the correlations between candidate samples online. Moreover, it does not set out to minimize the uncertainty from the learned GP model over untested measurements. Thus, large prediction errors could occur when the training data is constrained with few pilot surveys. Therefore, we propose to use information theoretic criteria with submodular framework that could characterize the correlations among the candidate samples online and hence improve the prediction accuracy from the learned GP model with sub-optimality bounds after selecting each new untested sample.

III. PROBLEM STATEMENT

In general, consider the entire sequenced set of n sensor measurements $\mathcal{V} = \{v_1, \dots, v_n\}$ where each measurement $\mathbf{v}_t = \{\text{locations, humidity, light, ...}\} \in \mathbb{R}^D$ for $t = 1, \dots, n$ corresponds to an environmental feature vector in D dimension feature space. There exists an unknown utility function $f: \mathbb{R}^D \rightarrow \mathbb{R}^+$ that maps the measurement input v_t with observable features to the scalar value of objective feature $f(v_t)$, in our testing case the temperature $y_t \in \mathbb{R}$. On each arrival of the measurement v_t , the robot will make an irrevocable decision¹ of whether to sample this measurement

¹We assume that revoking a decision requires a cost that is prohibitively large. For example, a robot collecting physical samples will have a large cost to clean the sample containers to prevent possible cross-contamination.

or not. Due to a budget limit, the robot will only be able to sample k measurements in total ($k < n$). After sampling k measurements, the robot will observe the actual utilities (the temperature reading in our case) of the measurements $\{y^i\}_{i=1}^k$ and use GP regression to learn the underlying mapping f , assuming the joint distribution of the observed offline temperature readings is Gaussian [6]. Hence, the goal is to select k best measurements under the stream-based setting for learning the mapping that minimizes the prediction error on new measurements.

A. Probabilistic Model

It has been shown that many environmental variables such as the temperature have a (multivariate) Gaussian joint distribution [6], [16], and hence we use the GP regression to learn the mapping from an environmental measurement $\mathbf{v}_t \in \mathbb{R}^D$ to the desired objective feature such as temperature, namely, compute the conditional posterior mean and variance for that objective variable as in [12]. Note that the mean function is assumed to be zero without loss of generality. Defining the training set of feature vectors as $\mathcal{A} \subset \mathcal{V}$, the conditional posterior mean $\mu_{\mathbf{v}_t|\mathcal{A}}$ in (1) is hence considered as the normalized predicted temperature for unobserved test measurement \mathbf{v}_t .

$$\begin{aligned} \mu_{\mathbf{v}_t|\mathcal{A}} &= \mathbf{k}(\mathbf{v}_t)^T (\mathbf{K}_{\mathcal{A}} + \sigma^2 \mathbf{I})^{-1} \mathbf{y} \\ \sigma_{\mathbf{v}_t|\mathcal{A}}^2 &= k(\mathbf{v}_t, \mathbf{v}_t) - \mathbf{k}(\mathbf{v}_t)^T (\mathbf{K}_{\mathcal{A}} + \sigma^2 \mathbf{I})^{-1} \mathbf{k}(\mathbf{v}_t) \end{aligned} \quad (1)$$

where $\mathbf{y} = [y^1, \dots, y^{|\mathcal{A}|}]^T$ is the noisy observation set for each training measurement $\mathbf{v}^i \in \mathcal{A}$ with $y^i = f(v^i) + \epsilon$ and Gaussian noise term $\epsilon \sim N(0, \sigma'^2)$. σ^2 is the measurement noise. $\mathbf{k}(\mathbf{v}_t) = [k(\mathbf{v}^1, \mathbf{v}_t), \dots, k(\mathbf{v}^{|\mathcal{A}|}, \mathbf{v}_t)]^T$ where $k(\mathbf{v}, \mathbf{v}')$ is the kernel function that captures the correlation between two measurement vectors. $\mathbf{K}_{\mathcal{A}}$ is the positive definite symmetric kernel matrix $[k(\mathbf{v}, \mathbf{v}')]_{\mathbf{v}, \mathbf{v}' \in \mathcal{A}}$. In particular, we use the following squared-exponential kernel function to specify the inter-sample correlation.

$$k(\mathbf{v}, \mathbf{v}') = \sigma_{SE}^2 e^{-\frac{(\mathbf{v}-\mathbf{v}')^T(\mathbf{v}-\mathbf{v}')}{2s^2}} \quad (2)$$

where the hyper-parameters s and σ_{SE} are length-scale and scale factor respectively that can be learned by optimizing the log marginal likelihood regarding the two parameters as done in [11].

In our case the actual value of the objective feature for the received measurements online cannot be observed until an offline analysis is processed, which makes the prediction $\mu_{\mathbf{v}_t|\mathcal{A}}$ unavailable for the online decision making due to unknown \mathbf{y} . To that end, we need metrics that capture the correlations between observed and untested measurements so as to select the most representative set of k measurements that reduce the uncertainty for prediction on untested measurements.

B. Objective Function

First, we consider the notion of differential entropy that is often used in spatial statistic optimization problems such as *experimental design* to specify the informativeness of a new unobserved sample given a set of previously selected unobserved measurements. Formally, for any new unobserved

measurement \mathbf{v}_t and existing training set of samples \mathcal{A} , the entropy of \mathbf{v}_t conditioned on set \mathcal{A} is defined as follows [6].

$$H(\mathbf{v}_t|\mathcal{A}) = \frac{1}{2} \log(2\pi e \sigma_{\mathbf{v}_t|\mathcal{A}}^2) \quad (3)$$

where $\sigma_{\mathbf{v}_t|\mathcal{A}}^2$ is computed in (1). While the differential entropy provides a good way to imply the reduction of uncertainty by the new unobserved measurement, it only concerns the entropy of selected measurements instead of the overall quality over the sampling space. To that end, an improved information theoretic criterion such as the mutual information [6], [12] is used here to indicate the reduction of uncertainty over the rest of the sampling space. Given the entire sampling set \mathcal{V} , the mutual information gain between any subset $\mathcal{A} \subset \mathcal{V}$ and the rest set $\mathcal{V} \setminus \mathcal{A}$ is defined as follows.

$$I(\mathcal{A}; \mathcal{V} \setminus \mathcal{A}) = H(\mathcal{V} \setminus \mathcal{A}) - H(\mathcal{V} \setminus \mathcal{A}|\mathcal{A}) \quad (4)$$

where $H(\mathcal{V} \setminus \mathcal{A})$ is the entropy of the rest sampling set that can be computed by chain rule [6]. Hence, our objective to select the best subset \mathcal{A}^* of k measurements is formally defined as follows.

$$\mathcal{A}^* = \arg \max_{\mathcal{A} \subset \mathcal{V}: |\mathcal{A}|=k} I(\mathcal{A}; \mathcal{V} \setminus \mathcal{A}) \quad (5)$$

Note that although the optimization criterion is similar to the *experimental design* problem [6], we do not assume that the knowledge of the entire sampling set \mathcal{V} is available in advance. In the stream-based setting as described in Fig. 1, the decision of sampling the current received measurement \mathbf{v}_t or not should be made before obtaining the information of the next measurement \mathbf{v}_{t+1} . Therefore, a stream-based optimal stopping strategy is needed to solve the combinatorial optimization problem (5).

IV. ONLINE SAMPLING USING STREAM-BASED SUBMODULAR SECRETARY ALGORITHM

The optimization problem (5) has been proven to be NP-hard [17]. Therefore, suboptimal solutions are acceptable. Given the additional constraint of irrevocable sampling decisions, we propose to combine optimal stopping theory and submodularity of the objective function (5) to *obtain a sub-optimal online decision making policy with provable bounds*.

A. Classical secretary algorithm

In the classical secretary problem, a company interviews n applicants sequentially without knowing their quality ranking at time of arrival, (and hence assumes that the arrival order is independently and identically distributed) with the goal to hire the person of the highest quality. The employer should make the irrevocable decision to hire or not right after interviewing each applicant. The solution is as follows: applicants are interviewed without hiring until the number of interviewed applicants reaches the cutoff point n/e . Then any applicant better than the best applicant until the last candidate will be hired. The optimal cutoff is n/e for a large n and the probability of choosing the best candidate is $1/e$ [13].

Consider the extension of the secretary problem to the multi-choice case where the goal is to select top k candidates from the stream of n candidates in total. Prior work [15] presented submodular secretary algorithms to select top k candidates given the submodularity of the monotone evaluation function. Such algorithms are particularly suitable for our problem since they do not suffer from highly correlated neighboring samples, as in other multi-choice secretary algorithms [18].

B. Submodularity Analysis

We prove that our objective function (5) is submodular.

Definition 1 (Submodularity [19]): Let \mathcal{V} be a finite set. A function $f: 2^{\mathcal{V}} \rightarrow \mathbb{R}$ is *submodular* if for all sets S and T with $S \subseteq T \subseteq \mathcal{V}$, the following is satisfied.

$$\forall v \notin T: f(S \cup \{v\}) - f(S) \geq f(T \cup \{v\}) - f(T) \quad (6)$$

Lemma 1: The set function of mutual information gain $\mathcal{A} \rightarrow I(\mathcal{A}; \mathcal{V} \setminus \mathcal{A})$ in (5) is submodular.

Proof: Let $\mathcal{A} \subseteq \mathcal{A}' \subseteq \mathcal{V}$ and $\mathbf{v} \in \mathcal{V} \setminus \mathcal{A}'$, which implies $\mathbf{v} \in \mathcal{V} \setminus \mathcal{A}$. For simplification we use $MI(\cdot)$ to denote $I(\cdot; \mathcal{V} \setminus \cdot)$. Then based on (4) we have

$$\begin{aligned} \Delta &= \{MI(\mathcal{A}' \cup \mathbf{v}) - MI(\mathcal{A}')\} - \{MI(\mathcal{A} \cup \mathbf{v}) - MI(\mathcal{A})\} \\ &= \{H(\mathbf{v}|\mathcal{A}') - H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A}' \cup \mathbf{v}))\} - \{H(\mathbf{v}|\mathcal{A}) - H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A} \cup \mathbf{v}))\} \\ &= \{H(\mathbf{v}|\mathcal{A}') - H(\mathbf{v}|\mathcal{A})\} + \{H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A} \cup \mathbf{v})) - H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A}' \cup \mathbf{v}))\} \end{aligned} \quad (7)$$

As more information indicates non-increasing entropy [6], [20], we have $H(\mathbf{v}|\mathcal{A}') \leq H(\mathbf{v}|\mathcal{A})$ and $H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A} \cup \mathbf{v})) \leq H(\mathbf{v}|\mathcal{V} \setminus (\mathcal{A}' \cup \mathbf{v}))$. Hence $\Delta \leq 0$ which concludes the proof. ■

[6] has proven that the mutual information gain is *approximately monotone* given fine discretization of the GP. Then together with *Lemma 1*, under the pool-based setting where \mathcal{V} is available in advance, we have the sub-optimality bound of $(1 - 1/e)$ using simple greedy algorithm that iteratively selects a candidate at each round t so that $\mathbf{v}_t = \arg \max MI(\mathcal{A}_{t-1} \cup \mathbf{v}_t)$ until $|\mathcal{A}_t| = k$.

As the objective function is submodular, next we present the adapted stream-based submodular sampling algorithm with performance bound.

C. The Stream-based Submodular Sampling Algorithm

We employ the general stream-based submodular secretary algorithms from [8], [15] and modify them in our problem considering that our objective function of mutual information gain $MI(\cdot)$ is monotone and submodular. The pseudocode is given in Algorithm 1 with inputs as sequenced set \mathcal{V} , sampled measurement set \mathcal{A} and cardinality constraint k .

Similar to [8], [15], Algorithm 1 first divides the measurement sequence $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ into k equally-sized selection windows defined by $S_l = \{\mathbf{v}_t \in \mathcal{V} | (l-1)n/k < t \leq ln/k\}$ for $l = 1, \dots, k$ (line 3).² In each segment (line 4–13) the classical secretary algorithm is applied to select one

²Without loss of generality we assume n is a multiple of k , or we can virtually fill some dummy measurements with zero incremental mutual information gain into the smaller segments to make them equivalent in length.

Algorithm 1 Stream-based Submodular Secretary Algorithm

```
1: procedure STREAMSUBMODULAR( $\mathcal{V}, \mathcal{A}, k$ )
2:    $\mathcal{A} \leftarrow \emptyset, a \leftarrow 0$ 
3:   for each segment  $S_l = \{\mathbf{v}_t \in \mathcal{V} | (l-1)n/k < t \leq ln/k\}$ 
4:     do
5:       for each measurement  $\mathbf{v}_t \in S_l$  do
6:         if  $(l-1)n/k < t < (l-1)n/k + n/(ke)$  then
7:           if  $MI(\mathcal{A} \cup \mathbf{v}_t) > a$  then
8:              $a \leftarrow MI(\mathcal{A} \cup \mathbf{v}_t)$ 
9:           end if
10:          else if  $MI(\mathcal{A} \cup \mathbf{v}_t) > a$  or  $\mathbf{v}_t$  the end in  $S_l$  then
11:             $\mathcal{A} \leftarrow \mathcal{A} \cup \mathbf{v}_t$ 
12:             $a \leftarrow 0$  and continue to the next segment  $S_{l+1}$ 
13:          end if
14:        end for
15:      end for
16:    end procedure
```

measurement with maximum mutual information gain on the existing selected sample set \mathcal{A} . Namely, during each observation phase (line 5–8) the robot will not sample the first nk/e new measurements while recording the maximum mutual information gain. During the following sampling phase (line 9–12), the incoming measurement will be sampled as long as (i) its mutual information gain is larger than the largest recorded value stored from the previous observation phase, or (ii) the current measurement is the end of the current segment S_l . Once the measurement is sampled, the robot will reject all the rest of the measurements in the current segment S_l and move on to the next segment S_{l+1} (line 11). Different from the algorithm design in [15] that selects at *most one* sample per segment, the Algorithm 1 ensures the selection of *exactly one* one measurement at each segment.

Performance Bound: Since we use the same segment partition assuming the identically distributed arrival order for the measurements as in [15] and the mutual information gain is proved to be submodular in *Lemma 1*, our Algorithm 1 shares the same sub-optimality bound $(1 - 1/e)/7$ as [15].

V. RESULTS

In this section, we present several simulation results on the benchmark dataset from Intel Berkeley Lab [9]. The dataset contains sensory data collected from 54 sensors in an office area between Feb 28th and Apr 5th, 2004. The data includes time-stamped readings such as sensor 2D locations, temperature, humidity, light and voltage. We want to predict temperature or light. The attributes in the dataset are divided into available training features and objective feature (temperature or light) to test the prediction performance. As described in Fig. 1, we assume that the robot starts from Sensor 1 and sequentially visits all the 54 sensors. Within the stream of 54 measurements, the robot makes k irrevocable sampling decisions. We compare the result of our approach described in Algorithm 1 to other competitive algorithms.

A. An Illustrative Example

First we consider an illustrative example where the robot collects 18 best measurements at one run using our stream-

based submodular secretary algorithm (Algorithm 1) to predict the temperature distribution of all points in the map. Since data on light, humidity and other variables are sensed only at the 54 sensor locations, we use only the 2D location information of the 18 sensors to learn the mapping to the temperature $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. For training purposes, we also use additional 5 randomly selected sensor data of location and temperature. After taking the 23 samples with location and temperature data in total as training measurements, the mapping from 2D location to temperature is learned to predict the temperature distribution over the whole map. Fig. 2a shows the ground-truth heatmap of temperature data in degrees Celsius fitted from sensory readings on the discrete locations of stationary sensors. Fig. 2b shows the predicted temperature distribution using only 18 selected measurements obtained from Algorithm 1 as well as the first five available measurements as randomly selected prior training data. With only data on location and from about half of the available sensors, our algorithm already shows good prediction performance compared with the one using the entire data stream.

To explicitly demonstrate the prediction performance compared to the ground truth for temperature estimation in the discrete locations of the non-selected sensors where data on light, humidity and other variables are also available, we run another trial using our algorithm with these additional dimensions of information. The inferred temperature distributions with variance over the discrete sensor locations is shown in Fig. 3a (with three prior measurements only) and Fig. 3b (after sampling another 20 selected measurements). It is straightforward to see that with our algorithm the posterior mean of predicted temperature over each sensor location is close to ground truth temperature reading using only 20 newly sampled data as the training set. The prediction variance is also largely reduced compared to the model from the three prior training measurements.

B. Numerical Results

To compare the algorithm performance with other similar works, we conducted experiments on 40 different streams of 54 sensory data from the Intel Berkeley data set randomly collected from Feb 28 to Mar 1, 2004 [9]. We use both mutual information gain (4) and Root Mean Square (RMS) error between predicted posterior mean and actual value for the objective feature to quantify the algorithm performance. On each experiment, different predefined number of selected samples are collected as different training datasets and the testing data is the sensor measurements recorded in the 54 sensor locations within the same time period on the following two days (Mar 2–3, 2004). In particular, we tested different algorithms for two different tasks: the temperature prediction and light prediction in the locations of the non-sampled sensors, using available measurements on sensor locations, humidity, voltage as well as temperature (when predicting light) or light (when predicting temperature).

We compare five representative online sampling algorithms (i.e., *Uniform*, *Multi-choice secretary* [7], *Stream GP-*

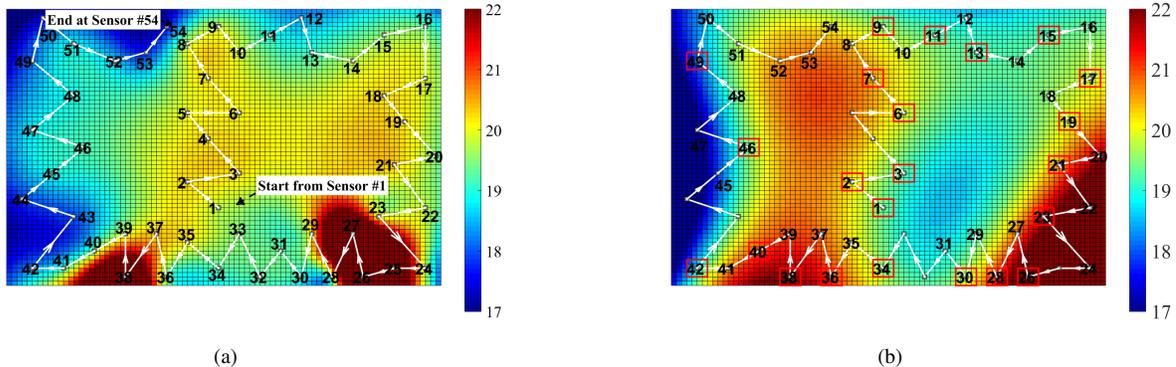


Fig. 2: An example of the temperature distribution modelling by sampling the best (according to Algorithm 1) 18 sensor location only information during a trip of the robot from sensor 1 to 54 (sequenced by the white arrows). (a) The ground-truth heat map of temperature in degrees Celsius with 47 available sensors (the readings from Sensor 4, 5, 25, 32, 33, 43, and 44 are not available in the dataset). (b) The modelled temperature distribution resulting from the best 18 available sensors selected plus first 5 available sensors that give location and temperature as training measurement set (both marked by the red box) by Algorithm 1.

UCB [1], *Stream Entropy*, *Stream Mean*) with our *Stream Submodular* secretary algorithm (Algorithm 1). The corresponding results are shown in Fig. 4 and Fig. 5 for inferring temperature and light, respectively. For the *Uniform* algorithm the robot uniformly selects the k measurements from the sequenced sensory data set regardless of value of observable features in each measurement. The *Multi-choice Secretary* algorithm [7] is a natural extension to the classical secretary algorithm. During the observation phase the first n/e measurements are passed without being selected and the best k measurements' incremental mutual information gain during this phase are recorded and ranked. In the following selection phase for the rest $n - n/e$ sequenced measurements, any new measurement will be selected as long as its incremental mutual information gain is larger than any one from the k records. The *Stream GP-UCB* [1] employed a similar multi-segment submodular algorithm as ours, but the sample selection criterion is based on the GP-UCB [12], which seeks for samples with maximum weighted sum of posterior mean and variance from (1) using pilot surveys. The *Stream Entropy* algorithm shares the same framework as our Algorithm 1, while using the entropy criterion (3) to select the best measurements. The *Stream Mean* also has the same framework as Algorithm 1 but with a different selection criterion that picks up measurements with highest posterior mean from (1). In Fig. 4 and Fig. 5, it is shown that our *Stream submodular* algorithm outperforms all the others in both prediction tasks given the resulting RMS error that directly indicates the prediction accuracy.

VI. CONCLUSION

We presented an online budget-restricted sampling algorithm under the stream-based setting. By using the mutual information gain as the selection criterion and exploiting the submodularity of the criterion, we improved the prediction accuracy of the learned GP model from budgeted sampled measurements and ensure the sub-optimality bound. Numerical simulations were performed on a real-world dataset to compare our stream-based submodular algorithm to other representative works in two different prediction tasks. The

results validated the effectiveness of the proposed algorithm.

Our current problem setting does not consider the actual objective feature value of those measurements (say temperature in our case) online until GP modelling. In the future we will investigate the situation that allows immediate realization of the objective feature value after sampling each measurement, so that we can better reduce the prediction uncertainty or improve the utility of the selected measurements in an adaptive manner. To further test and refine our algorithm we also plan to conduct field experiment with our sampling algorithm on data samples of sorghum varieties to predict the varieties' yields and susceptibility to disease.

REFERENCES

- [1] J. Das, F. Py, J. B. Harvey, J. P. Ryan, A. Gellene, R. Graham, D. A. Caron, K. Rajan, and G. S. Sukhatme, "Data-driven robotic sampling for marine ecosystem monitoring," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1435–1452, 2015.
- [2] T. Mueller-Sim, M. Jenkins, J. Abel, and G. Kantor, "The robotanist: A ground-based agricultural robot for high-throughput crop phenotyping," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 3634–3639.
- [3] A. Singh, F. Ramos, H. D. Whyte, and W. J. Kaiser, "Modeling and decision making in spatio-temporal processes for environmental surveillance," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 5490–5497.
- [4] W. Luo, S. S. Khatib, S. Nagavalli, N. Chakraborty, and K. Sycara, "Distributed knowledge leader selection for multi-robot environmental sampling under bandwidth constraints," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 5751–5757.
- [5] K. Chaloner and I. Verdinelli, "Bayesian experimental design: A review," *Statistical Science*, pp. 273–304, 1995.
- [6] 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 (JMLR)*, vol. 9, pp. 235–284, February 2008.
- [7] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg, "A knapsack secretary problem with applications," in *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Springer, 2007, pp. 16–28.
- [8] K. Fujii and H. Kashima, "Budgeted stream-based active learning via adaptive submodular maximization," in *Advances in Neural Information Processing Systems*, 2016, pp. 514–522.
- [9] (2004) Intel lab data. [Online]. Available: <http://db.csail.mit.edu/labdata/labdata.html>
- [10] G. A. Hollinger and G. S. Sukhatme, "Sampling-based robotic information gathering algorithms," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1271–1287, 2014.

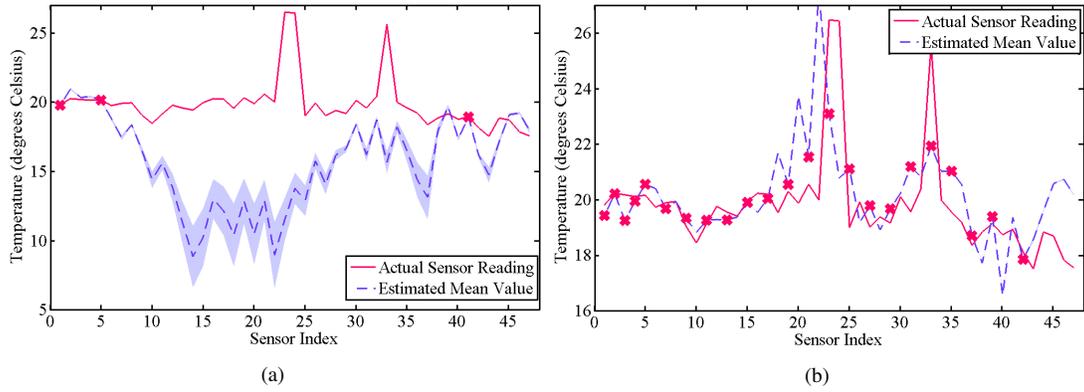


Fig. 3: Comparisons of temperature estimation before and after the online sampling. (a) The temperature estimation at 47 sensor locations from the pilot study with first three available sensor readings (marked by red forks). The blue area indicates the posterior variance over the respective sensor location. (b) The temperature estimation after sampling the best 20 sensors information and the first three (both marked by red forks).

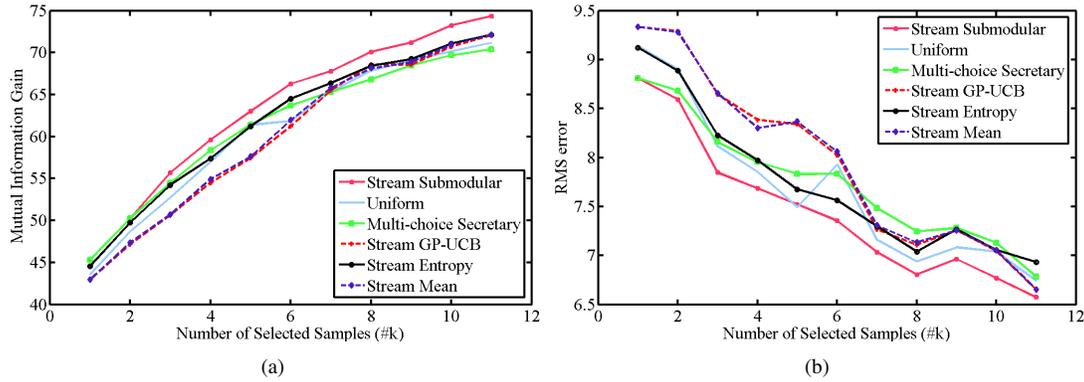


Fig. 4: Comparisons of the mutual information gain and RMS error vs. predefined number of selected samples by different algorithms and criteria (temperature prediction).

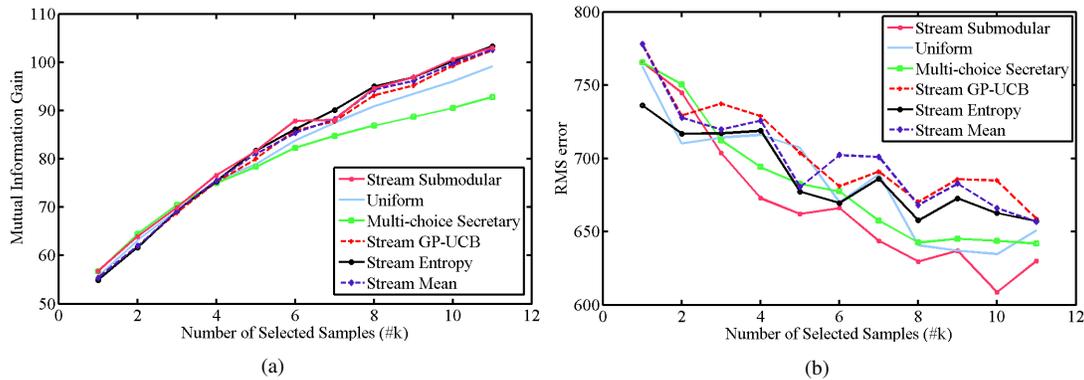


Fig. 5: Comparisons of the mutual information gain and RMS error vs. predefined number of selected samples by the different algorithms and criteria (light prediction).

[11] C. E. Rasmussen and C. K. Williams, *Gaussian processes for machine learning*. MIT press Cambridge, 2006, vol. 1.

[12] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, "Information-theoretic regret bounds for gaussian process optimization in the bandit setting," *IEEE Transactions on Information Theory*, vol. 58, no. 5, pp. 3250–3265, 2012.

[13] P. Freeman, "The secretary problem and its extensions: A review," *International Statistical Review*, vol. 51, pp. 189–206, 1983.

[14] T. S. Ferguson, "Who solved the secretary problem?" *Statistical science*, pp. 282–289, 1989.

[15] M. Bateni, M. Hajiaghayi, and M. Zadimoghaddam, "Submodular secretary problem and extensions," *ACM Transactions on Algorithms (TALG)*, vol. 9, no. 4, p. 32, 2013.

[16] A. Deshpande, C. Guestrin, S. R. Madden, J. M. Hellerstein, and W. Hong, "Model-driven data acquisition in sensor networks," in *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*. VLDB Endowment, 2004, pp. 588–599.

[17] C.-W. Ko, J. Lee, and M. Queyranne, "An exact algorithm for maximum entropy sampling," *Operations Research*, vol. 43, no. 4, pp. 684–691, 1995.

[18] Y. Girdhar and G. Dudek, "Optimal online data sampling or how to hire the best secretaries," in *Computer and Robot Vision, 2009. CRV'09. Canadian Conference on*. IEEE, 2009, pp. 292–298.

[19] S. Fujishige, *Submodular functions and optimization*. Elsevier, 2005, vol. 58.

[20] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.