Selecting Most Informative Contributors with Unknown Costs for Budgeted Crowdsensing

Shuo Yang†, Fan Wu†, Shaojie Tang‡, Tie Luo§, Xiaofeng Gao†, Linghe Kong†, Guihai Chen†
†Shanghai Key Laboratory of Scalable Computing and Systems, Shanghai Jiao Tong University, China
‡Department of Information Systems, University of Texas at Dallas, USA
§Institute for Infocomm Research, A*STAR, Singapore
†{wnmmxy, wu-fan, gao-xf, linghe.kong, gchen}@sjtu.edu.cn; ‡tangshaojie@gmail.com;§luot@i2r.a-star.edu.sg

***Abstract*—Mobile crowdsensing has become a novel and promising paradigm in collecting environmental data. A critical problem in improving the QoS of crowdsensing is to decide which users to select to perform sensing tasks, in order to obtain the most informative data, while maintaining the total sensing costs below a given budget. The key challenges lie in (i) finding an effective measure of the informativeness of users' data, (ii) learning users' sensing costs which are unknown a priori, and (iii) designing efficient user selection algorithms that achieve low-regret guarantees. In this paper, we build Gaussian Processes (G-Ps) to model spatial locations, and provide a mutual information-based criteria to characterize users' informativeness. To tackle the second and third challenges, we model the problem as a budgeted multi-armed bandit (MAB) problem based on stochastic assumptions, and propose an algorithm with theoretically proven low-regret guarantee. Our theoretical analysis and evaluation results both demonstrate that our algorithm can efficiently select most informative users under stringent constraints.**

I. Introduction

One of the important QoS problems that have not been carefully addressed in mobile crowdsensing is how to reduce the redundancy of collected information. Specifically, we intend to select users whose data are most informative about the monitored environment. On the one hand, environmental conditions of two nearby locations tend to be similar [1]. If the sensing data of some specific locations have already been collected, it is of less necessity to collect nearby information. On the other hand, in many cases, environmental conditions have certain spatial correlations [2], which allow us to model existing observations and make predictions for unobserved ones. Thus, it is of great significance to select users who are most informative about unobserved locations, especially when there is only limited budget to pay the users.

To motivate users' participation, the crowdsensing platform usually needs to pay a certain amount of reward to selected users. Given a limited budget, a natural idea is to make use of the budget feasible mechanisms [3], [4]. However, directly applying these approaches to crowdsensing scenarios incurs at least two practical limitations.

First, letting users determine their own deserved payments has critical drawbacks. Existing works on crowdsensing (*e.g.*, [4], [5]) are usually based on a reverse-auction model, where each user is assumed to be aware of her sensing cost and submits her cost as a reverse price, or bid, for performing a requested task. However, this assumption may not be practical in reality, since users' contributions in a crowdsensing task are not purely financial, but involve hardware consumptions and their data qualities. Quantifying the consumption of users' hardware resources requires complicated technical skills, and hence is infeasible for most users. Furthermore, the payment that each user deserves should be related to her data quality [6]. Intuitively, users with high data qualities deserve higher payments, and vice versa. Thus, instead of passively listening to users' bids, we prefer a more "aggressive" crowdsensing platform that actively quantifies users' contributions and determines their deserved payments in a proper manner.

Second, even though the users' hardware consumptions can be quantified, such information is unknown a priori. In other words, the platform can access to a user's sensing cost only after the user has finished the task and submitted her data. The unavailability of prior knowledge on users' sensing costs brings us another challenge in user selection. Suppose we have already found some cheap and informative users, we need to decide whether to keep selecting these users to maintain good performance, or to select other users in hope of finding even cheaper and more informative ones. This is an instance of the *exploration* and *exploitation* dilemma in reinforcement learning [7], where an agent needs to decide whether to explore new information about the effectiveness of an action, or to exploit the action that is already known to be effective.

In this paper, we address the problem of selecting the most informative users, while the platform has no prior knowledge of users' sensing costs and the total payments are limited by a fixed budget. We model the spatial environment using the Gaussian Processes (GPs), and adopt a mutual information-based criteria to quantify the informativeness of users. Next, we consider an unrealistic but instructive scenario where the platform has full knowledge of users' costs. To tackle the NP-hardness of the budgeted maximization of submodular functions, we propose an efficient multi-rounded algorithm that achieves $(1-1/e)/2$ approximation ratio. Then, we

This work was supported in part by the State Key Development Program for Basic Research of China (973 project 2012CB316201), in part by China NSF grant 61422208, 61472252, 61272443, 61133006 and 61303202, in part by Shanghai Science and Technology fund 15220721300, in part by CCF-Tencent Open Fund and Open Project of Baidu 181515P005267, in part by the Opening Project of Key Lab of Information Network Security of Ministry of Public Security C15602. The opinions, findings, conclusions, and recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies or the government.

F. Wu is the corresponding author.

consider a realistic scenario where the prior knowledge of users' costs is not available, and propose an efficient Budgeted Informativeness Maximization algorithm, namely BIM, to actively learn users' costs and decide which users to select. Our theoretical analysis shows that BIM achieves zero regret in an asymptotic case. We also evaluate BIM in various settings. Our evaluation results demonstrate good performance of BIM.

The rest of this paper is organized as follows. We first present a system overview in Section II. Then, we address the budgeted informativeness maximization with and without prior knowledge in Section III and Section IV, respectively. In Section V, we evaluate our algorithm and present the evaluation results. The related works are presented in Section VI. Finally, we conclude this paper in Section VII.

II. System Overview

A. Problem Statement

A typical crowdsensing architecture consists of three major components: service requesters, mobile device users, and a crowdsensing platform. After receiving location-based sensing requests from the service requesters, the platform releases specific sensing tasks to the mobile device users (we will refer as users for simplicity). Without loss of generality, we focus on one task only, *e.g.*, noise monitoring at a specified park. We assume that the task consists of T consecutive rounds and each round has a fixed duration. The specified sensing area is represented by a set of finite discrete locations $\mathcal{L} = \{l_1, l_2, \ldots, l_m\}$, *e.g.*, a grid discretization of $\mathbb{R}^2$.

Suppose there is a set $\mathcal{N} = \{1, 2, \ldots, n\}$ of users interested in the task. At the beginning of each round $t \in \{1, 2 \ldots, T\}$, the platform selects a subset of users $\mathcal{S}_t \subseteq \mathcal{N}$ to sense in this round. After round t, each selected user $i \in \mathcal{S}_t$ submits her sensing data and receives her payment $p_{i,t}$ of this round. Users' real-time locations are measured by the GPS modules of their smartphone and reported to the platform via WiFi or celluar network. During each round t, each selected user $i \in \mathcal{S}_t$ is required to stay at a same location, denoted by $l_{i,t}$.

The main objective of the crowdsensing campaign is to maximize the total informativeness within a given budget B. We let $F(\mathcal{S})$ denote how informative a set $\mathcal{S} \subseteq \mathcal{N}$ of users is. Our problem is formalized as follows:

$$\max_{\mathcal{S}_t \subseteq \mathcal{N}} \sum_{t=1}^{T} F(\mathcal{S}_t), \text{subject to} \sum_{t=1}^{T} \sum_{i \in \mathcal{S}_t} p_{i,t} \leq B. \tag{1}$$

B. Quantifying Informativeness

We adopt the mutual information criteria [8] to model users' informativeness. We associate a random variable X_{l_s} with each location $l_s \in \mathcal{L}$. For a subset $\mathcal{S} \subseteq \mathcal{N}$, let $\mathcal{L}_\mathcal{S} \subseteq \mathcal{L}$ denote the set of locations of users $\mathcal{S}$, then $X_{\mathcal{L}_\mathcal{S}}$ is the set of random variables associated with the locations $\mathcal{L}_\mathcal{S}$. To simplify notation, we write l_s instead of X_{l_s} and $L_\mathcal{S}$ instead of $X_{\mathcal{L}_\mathcal{S}}$.

Definition 1 (Mutual Information [8]). *The mutual information $I(X;Y)$ between two random variables X and Y are*

$$I(X;Y) = \int_{\mathcal{Y}} \int_{\mathcal{X}} p(x,y)\, log \frac{p(x,y)}{p(x)p(y)} dx\, dy. \tag{2}$$

Intuitively, mutual information measures the information that X and Y share, *i.e.*, how much knowing one reduces uncertainty about the other. For example, if X and Y are independent, then knowing X does not give any information about Y and vice versa, and thus their mutual information is zero. In monitoring spatial environment, suppose we have already deployed users $\mathcal{S} \subseteq \mathcal{N}$ to perform the sensing task, then $I(\mathcal{L}_\mathcal{S}; \mathcal{L} \backslash \mathcal{L}_\mathcal{S})$ represents how much knowledge observing $\mathcal{L}_\mathcal{S}$ gives about the unobserved locations $\mathcal{L} \backslash \mathcal{L}_\mathcal{S}$. According to Equation (2) and (3), we have

$$I(X;Y) = H(X) + H(Y) - H(X,Y), \tag{3}$$

where $H(X,Y) = -E[\log p(X,Y)]$ is the the joint entropy of X and Y. We define the informativeness function as:

$$F(\mathcal{S}) = I(\mathcal{L}_\mathcal{S}; \mathcal{L} \backslash \mathcal{L}_\mathcal{S}), \tag{4}$$

so that maximizing the informativeness of selected users is equivalent to finding the subset $\mathcal{S}$ of users such that the mutual information between $\mathcal{L}_\mathcal{S}$ and $\mathcal{L} \backslash \mathcal{L}_\mathcal{S}$ is maximized.

C. Multivariate Gaussian Distribution and Gaussian Process

In crowdsensing, it is often desirable not only to predict the values of environmental conditions, but also to estimate their probabilistic distributions. Following [2], we assume that the monitored phenomena follow a joint multivariate Gaussian distribution. Mathematically, the joint distribution of a set $X_\mathcal{L} = \{X_{l_1}, X_{l_2}, \ldots, X_{l_m}\}$ of random variables with m discrete locations is

$$P(X_\mathcal{L} = x_\mathcal{L}) = \frac{1}{(2\pi)^{m/2} |\Sigma|^{1/2}} e^{-\frac{1}{2}(x_\mathcal{L} - \mu)^T \Sigma^{-1} (x_\mathcal{L} - \mu)}, \tag{5}$$

where μ is the mean vector, Σ is the covariance matrix, and $|\Sigma|$ the determinant of Σ. We note that the conditional distribution of a joint Gaussian distribution is still Gaussian. If in some round t, we have observed the noise of locations $\mathcal{L}_t \subseteq \mathcal{L}$, then for some unobserved location $l_s \in \mathcal{L} \backslash \mathcal{L}_t$, its conditional mean $\mu_{l_s|\mathcal{L}_t}$ and conditional variance $\sigma^2_{l_s|\mathcal{L}_t}$ are given by:

$$\mu_{l_s|\mathcal{L}_t} = \mu_{l_s} + \Sigma_{l_s \mathcal{L}_t} \Sigma^{-1}_{\mathcal{L}_t \mathcal{L}_t} (x_{\mathcal{L}_t} - \mu_{\mathcal{L}_t}), \tag{6}$$

$$\sigma^2_{l_s|\mathcal{L}_t} = \sigma^2_{l_s} - \Sigma_{l_s \mathcal{L}_t} \Sigma^{-1}_{\mathcal{L}_t \mathcal{L}_t} \Sigma_{\mathcal{L}_t l_s}. \tag{7}$$

The $\Sigma_{l_s \mathcal{L}_t}$ is a vector of the covariance of l_s with all the variables in $\mathcal{L}_t$, and $\Sigma_{\mathcal{L}_t \mathcal{L}_t}$ is the covariance matrix where the entry (u, v) is the covariance of u and v.

In some cases, we are not only interested in specified locations, but also those unspecified locations. Gaussian process can be utilized to generalize multivariate Gaussian distribution to scenarios with infinite number of random variables. In our noise monitoring example, we can have infinite number of location indexes, *e.g.*, $\mathcal{L} \subseteq \mathbb{R}^2$, and each location l_s is associated with a random variable X_s. The Gaussian process is specified by a mean function $M(\cdot)$ and a symmetric positive definite kernel function $K(\cdot, \cdot)$. For each random variable X_u, $M(u)$ represents its mean, and for any two random variables X_u and X_v, their covariance is represented by $K(u, v)$.

Given a set $X_{\mathcal{L}_\mathcal{S}}$ of k random variables, if they follow a multivariate normal distribution with mean vector $\mu_\mathcal{S}$ and covariance matrix $\Sigma_{\mathcal{S}\mathcal{S}}$, then their (differential) entropy can be calculated by [8]:

$$H(X_{\mathcal{L}_\mathcal{S}}) = \frac{1}{2}\log((2\pi e)^k|\Sigma_{\mathcal{S}\mathcal{S}}|) \tag{8}$$

Based on Equation (4) and (8), we are now able to estimate the mutual information between locations of selected users and unobserved locations.

D. A Multi-Armed Bandit Problem

The void of prior knowledge of users' costs brings us the exploration and exploitation dilemma, originally proposed by Robbins [9] as a multi-armed bandit (MAB) problem. In a classical stochastic MAB framework [7], there is a slot machine with multiple non-identical arms, and pulling each arm generates a random reward according to some unknown distribution with unknown mean. A gambler must decide which arms to pull in sequence, with the objective of minimizing *regret*, *i.e.*, the difference between the optimal payoff with expert knowledge and the actual payoff.

In our problem, we model the set of mobile device users as n non-identical arms. Selecting user i in round t incurs a cost $c_{i,t}$, which follows some unknown distribution with unknown mean c_i. At the beginning of each round t, the crowdsensing platform (gambler) decides which users $\mathcal{S}_t \subseteq \mathcal{N}$ to perform the sensing task. After round t, the platform obtains each selected user i's real cost $c_{i,t}$ in round t, and updates i's expected cost. Each selected user $i \in \mathcal{S}_t$ will receive a payment $p_{i,t}$, which is proportional to her cost $c_{i,t}$. The exact calculation formula depends on specific crowdsensing scenarios, and is left for the platform to decide.

A plenty of MAB algorithms have been proposed based on different approaches (see a survey in [10]). However, our problem has key differences from them. First, most of the existing works assume that one or fixed number of arm(s) can be pulled at a time, while in our cases, the number of selected users in each round is uncertain. Though some work considers a combinatorial MAB approach (*e.g.*, [11]), it never takes the budget constraint into consideration. Second, the optimization functions of most previous works are either additive [12]–[17] or symmetric (the function output only depends on the cardinality of the input set) [18], but in our cases, the mutual information-based criteria is a submodular function, which is more challenging. Third, in contrast to many previous works on budgeted MAB (*e.g.*, [12], [14], [15]), where they consider maximizing unknown profits (with additive optimization function) with known costs, we have no prior knowledge of users' costs. Although [16]–[18] consider variable costs, their problem models are quite different from ours (see our technical report [19] for more detailed comparisons).

III. Budgeted Informativeness Maximization with Full Knowledge

In this section, we address the budgeted informativeness maximization problem with full knowledge of users' sensing costs. Although the full knowledge assumption is unrealistic, the algorithms proposed in this section are the building blocks of our subsequent designs without this assumption.

Algorithm 1: Single-Rounded Budgeted Informativeness Maximization with Full Knowledge

Input: The total budget B and users' costs $\{c_i\}$
Output: Selected user $\mathcal{S}$

Calculate each user's payment p_i based on c_i;
$i^* \leftarrow \underset{i\in\mathcal{N}, p_i \le B}{\text{argmax}}\ F(\{i\})$;
$\mathcal{S}' \leftarrow \{i^*\}$, $\mathcal{S}'' \leftarrow \emptyset$, $\mathcal{N}' \leftarrow \mathcal{N}$, $B' \leftarrow B$;
while $\mathcal{N}' \neq \emptyset$ **do**
  $i^* \leftarrow \underset{i\in\mathcal{N}'}{\text{argmax}} \frac{F(\mathcal{S}''\cup\{i\})-F(\mathcal{S}'')}{p_i}$;
  if $B' \ge p_{i^*}$ **then** $\mathcal{S}'' \leftarrow \mathcal{S}'' \cup \{i^*\}$, $B' \leftarrow B' - p_{i^*}$;
  $\mathcal{N}' \leftarrow \mathcal{N}' \backslash \{i^*\}$;
return $\underset{\mathcal{S}\in\{\mathcal{S}',\mathcal{S}''\}}{\text{argmax}}\ F(\mathcal{S})$;

A. Single-Rounded Budgeted Informativeness Maximization

We first present the definition of submodular functions, which have a natural diminishing return property, *i.e.*, the marginal gain when adding a single element to a input set decreases as the size of the input set increases.

Definition 2 (Submodular)**.** *The function $F : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ is submodular if $\forall \mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{N}, \forall i \in \mathcal{N}\backslash\mathcal{B}$,*

$$F(\mathcal{A}\cup\{i\}) - F(\mathcal{A}) \ge F(\mathcal{B}\cup\{i\}) - F(\mathcal{B}). \tag{9}$$

Although mutual information is not submodular in general, our informativeness characterization $F(\mathcal{L}_\mathcal{S}) = I(\mathcal{L}_\mathcal{S}; \mathcal{L}\backslash\mathcal{L}_\mathcal{S})$ is a non-negative submodular function. The function F, with a careful discretization over the monitored environmental area, is non-decreasing over $\mathcal{L}_\mathcal{N}$ [20].

Lemma 1. *The function $F(\mathcal{S}) = I(\mathcal{L}_\mathcal{S}; \mathcal{L}\backslash\mathcal{L}_\mathcal{S})$ is non-negative and submodular.*

Proof. Please refer to our technical report [19]. □

When the crowdsensing task has only a single round, all we need to do is to select a subset of users that maximize the informativeness function F under the budget constraint. The problem is NP-hard even with the full knowledge of users' costs [21]. A simple greedy algorithm that greedily selects the next user that maximizes the marginal informativeness gain/cost ratio before the budget drains out has unbounded approximation ratio. But with a slight modification [22], it can achieve $(1 - 1/e)/2$ approximation to the optimal selection. The algorithm is presented in Algorithm 1.

B. Multi-Rounded Budgeted Informativeness Maximization

In our cases, the multi-rounded user selection, formalized in Equation (1), is more challenging, since the total budget over the entire T rounds is limited. A careful design is required on how to allocate the total budget among the T rounds.

A simple equally shared scheme can be arbitrarily bad. For example, consider a two-rounded crowdsensing with only two users, denoted by n_1 and n_2. The costs of n_1 and n_2 are 2 and 1 per round respectively. Suppose these two users have fixed locations, and their informativeness is characterized as follows: $F(\{n_1\}) = o \gg 1$ and $F(\{n_2\}) = 1$. If the total budget is 2, then one optimal solution is to choose n_1 for

Algorithm 2: Multi-Rounded Budgeted Informativeness Maximization with Full Knowledge

Input: Budget B, the total round T, and users' costs $\{c_i\}$
Output: User selections $\{\phi_i\}, \forall i \in \mathcal{N}$

1 Calculate each user's payment p_i based on c_i;
2 $\langle i^*, \phi^*_{i^*} \rangle \leftarrow \underset{i\in\mathcal{N},\phi_i\in\Phi,|\phi_i|p_i\leq B}{\text{argmax}} G(\{\phi_i\})$;
3 $\mathcal{V}' \leftarrow \{\phi^*_{i^*}\}, \mathcal{V}'' \leftarrow \emptyset, \mathcal{N}' \leftarrow \mathcal{N}, B' \leftarrow B$;
4 **while** $\mathcal{N}' \neq \emptyset$ **do**
5   **foreach** $i \in \mathcal{N}'$ **do**
6   **if** $B' \geq p_i$ **then**
7       $\phi^*_i \leftarrow \underset{\phi_i\in\Phi,|\phi_i|p_i\leq B'}{\text{argmax}} G(\mathcal{V}'' \cup \{\phi_i\}) - G(\mathcal{V}'')$;
8   **else** $\mathcal{N}' \leftarrow \mathcal{N}' \backslash \{i\}$;
9   **if** $\mathcal{N}' \neq \emptyset$ **then**
10     $i^* \leftarrow \underset{i\in\mathcal{N}'}{\text{argmax}} \frac{G(\mathcal{V}''\cup\{\phi^*_i\})-G(\mathcal{V}'')}{|\phi^*_i|p_i}$;
11     $\mathcal{V}'' \leftarrow \mathcal{V}'' \cup \{\phi^*_{i^*}\}, \mathcal{N}' \leftarrow \mathcal{N}' \backslash \{i^*\}$;
12     $B' \leftarrow B' - |\phi^*_{i^*}|p_i$;
13 **return** $\underset{\mathcal{V}\in\{\mathcal{V}',\mathcal{V}''\}}{\text{argmax}} G(\mathcal{V})$;

the first round and neither for the second round. The optimal result is o. However, if we divide the budget equally into these two rounds, then the user selection ends up with n_2 in both rounds, with the total output being 2. The approximation ratio is $o/2$, which is unbounded and can be arbitrarily bad.

To tackle the NP-hardness of the multi-rounded budgeted informativeness maximization, we extend the single-rounded approximation algorithm to support multi-rounded user selection with $(1 - 1/e)/2$ approximation guarantee, when users' locations are fixed during the entire T rounds.

The algorithm is presented in Algorithm 2. Instead of trying to allocate the total budget among the T rounds, we consider a global coordination of user selections. For each user $i \in \mathcal{N}$, we define ϕ_i to be a T-ary vector of user i's selection instance, where $\phi_{i,t} \in \{0, 1\}$ denotes if user i has been selected in round t ("1" denotes selected and "0" otherwise). Let Φ denote the set of all possible permutations of the T-ary vector. We define $|\phi_i|$ to be how many times user i has been selected among T rounds, *i.e.*, $|\phi_i| = \sum_{i=1}^{T} \phi_{i,t}$. The set $\{\phi_1, \phi_2, \ldots, \phi_n\}$ is denoted by $\mathcal{V}$. The extended informativeness function $G : 2^{\mathcal{V}} \rightarrow \mathbb{R}$ is defined as follows:

$$G(\mathcal{V}) = \sum_{i=1}^{T} F(\mathcal{S}_t), \tag{10}$$

where $\mathcal{S}_t = \{i|\phi_{i,t} = 1, i \in \mathcal{N}, \phi_i \in \mathcal{V}\}$. We can see that G is a nondecreasing submodular function, since the addition of several monotone submodular functions is still monotone submodular. Now, our problem can be treated as selecting a set $\mathcal{V}'$ of users' selection instances that maximizes the submodular function G under the budget constraint B.

Our algorithm follows the similar design rationale of the single-rounded approximation algorithm. For each user i, we first calculate her best affordable selection instance $\phi^*_i \in \Phi$ that maximizes the marginal informativeness gain under current user selection $\mathcal{V}''$ (Line 7), *i.e.*, $G(\mathcal{V}'' \cup \{\phi_i\}) - G(\mathcal{V}'')$. Then, we greedily select the best user i^* that maximizes the marginal informativeness gain of her best instance over the instance's cost (Line 10). We note that for each user i, her current best selection instance ϕ^* (Line 7) can be calculated efficiently by sequentially activating $\phi_{i,t} = 1$ for some round t that maximizes the the marginal informativeness gain before the budget drains out. Afterwards, the performance of this greedy policy $G(\mathcal{V}'')$ is then compared with the best performance that selecting only one user can achieve, *i.e.*, $G(\mathcal{V}')$, and the better user selection is returned.

Theorem 1. *The proposed multi-rounded budgeted informativeness maximization algorithm can achieve an approximation ratio of* $\frac{1-1/e}{2}$*, with polynomial time complexity.*

Proof. The proof is similar to the proof of the single-rounded algorithm [22]. We omit the proof due to space limitation. □

IV. Budgeted Informativeness Maximization without Prior Knowledge

Under the cases where the knowledge of users' costs are unknown in advance, we need to learn users costs. We propose a simple but effective method to balance the exploration and the exploitation. The intuitive idea of our algorithm is that we allocate a portion of the total budget $B' = \epsilon B, \epsilon \in (0, 1)$ to learn users' costs, and the remaining budget will be used to choose the best users.

Our algorithm, namely BIM, is shown in Algorithm 3. It consists of an exploration phase and an exploitation phase. In the exploration phase, we select all the users in each round, so as to collect information about users' sensing costs. After each round t, the platform quantifies each user i's sensing cost $c_{i,t}$, and calculate each user's payment $p_{i,t}$ according to her cost. Suppose the maximum payment the platform is willing to pay each user per round is p_{max}, which is high enough to cover any user's sensing cost in a single round. The p_{max} is only used to estimate if the remaining budget of B' can afford another exploration round (Line 2). After the exploration phase, we calculate the estimate of each user i's sensing cost $\{\hat{c}_i\}$ (Line 8), and recycle the remaining budget of the exploration phase (Line 9). In the exploitation phase, we treat $\hat{c}_i$ as each user i's true sensing cost, and apply Algorithm 2 to determine the user selections of the following rounds.

To analyze the performance of BIM, we need to calculate its *regret*, which is the difference between its obtained informativeness and the optimum. However, the optimal informativeness cannot be feasibly achieved, due to the NP-hardness of the submodular maximization problem. To address this infeasibility in regret analysis, we adopt the concept of α-*approximation regret*.

Definition 3 (α-approximation regret [11], [21])**.** *The α-approximation regret of a sequence of user selection $\{\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_T\}$ is*

$$R_\alpha = \alpha \cdot \sum_{t=1}^{T} F(\mathcal{S}^*_t) - \sum_{t=1}^{T} F(\mathcal{S}_t), \tag{11}$$

where $\{\mathcal{S}^*_1, \mathcal{S}^*_2, \ldots, \mathcal{S}^*_T\}$ is the optimal user selection sequence. The average α-approximation regret is R_α/T.

Algorithm 3: BIM

Input: The total budget B, payment upper bound p_{max}, a constant ϵ
Output: Selected user $\{\mathcal{S}_1, \ldots, \mathcal{S}_T\}$
$B' \leftarrow \epsilon B$, $B'' \leftarrow (1-\epsilon)B$, $T' \leftarrow 0$;
// Exploration Phase
while $B' \geq np_{max}$ *and* $T' < T$ **do**
  $T' \leftarrow T' + 1$;
  $\mathcal{S}_{T'} \leftarrow \mathcal{N}$;
  Collect users data, and calculate $c_{i,T'}$ and $p_{i,T'}, \forall i \in \mathcal{N}$;
  $B' \leftarrow B' - \sum_{i \in \mathcal{S}_{T'}} p_{i,T'}$;
// Exploitation Phase
$\hat{c}_i \leftarrow \sum_{t=1}^{T'} c_{i,t}/T', \forall i \in \mathcal{N}$;
$B'' \leftarrow B'' + B'$;
$\mathcal{V} \leftarrow \text{Alg}_2(B'', T - T', \{\hat{c}_i\})$;
Calculate $\mathcal{S}_t, \forall t \geq T' + 1$ based on $\mathcal{V}$;
return $\{\mathcal{S}_1, \ldots, \mathcal{S}_T\}$;

The intuitive meaning of α-approximation regret is that our regret metric does not compare against the optimal user selection sequence, but against an α-approximation oracle. Based on the above definition, we can analyze the $\frac{1-1/e}{2}$-approximation regret of BIM, *i.e.*, to compare the performance of BIM against Algorithm 2.

Theorem 2. *The $\frac{1-1/e}{2}$-approximation regret of BIM is* $(T - \lfloor \frac{B}{np_{max}} \rfloor)F(\mathcal{N})$.

Proof. Please refer to our technical report [19]. □

V. Evaluation

In this section, we evaluate the performance of BIM by comparing it with two benchmarks. One of the benchmarks is the full-knowledged greedy algorithm. The other is a random scheme that randomly selects users until the budget drains out.

A. Simulation Setup

We consider a squared sensing area of $1km \times 1km$. The area is discretized into 400 locations, and each location is a $50m*50m$ grid. The mean of each user i's cost c_i is uniformly generated in (5,10), and i's cost in each round t follows a Gaussian distribution with mean c_i and variance 0.2, *i.e.*, $c_{i,t} \sim N(c_i, 0.2)$. Users' locations are randomly generated and fixed during the entire task. The Gaussian process is characterized by a classic Gaussian Kernel:

$$K(u,v) = \exp(-\frac{\| u,v \|^2}{h^2}), \tag{12}$$

where $\| u,v \|^2$ is the squared Euclidean distance of u and v, and h is a constant parameter. We set h=1000. The range of user number n and round number T vary from 5 to 100 with the increment of 5. The budget B varies from 1000 to 11000 with the increment of 1000. The default n, T, and B are 50, 20, and 5000 respectively. The constant ϵ is set to 0.5.

B. Evaluation Results

Fig. 1 varies n from 5 to 100 with fixed T and B. We can see that the total informativeness increases as n gets larger, and finally tends to reach a stable state, when the budget has

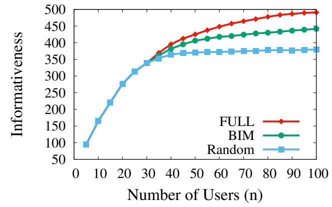

Fig. 1. Varying n

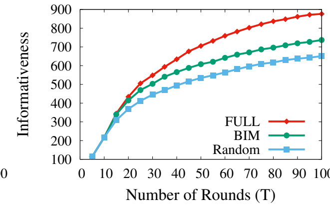

Fig. 2. Varying T

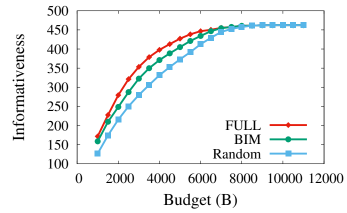

Fig. 3. Varying B

been fully used. It can also be observed that when n is small, *e.g.*, less than 30, all the three algorithms achieve the same informativeness. This is because the total budget can afford all the users to perform the sensing task in each round. When n gets larger, the performance differences of these algorithms start to reveal, *i.e.*, full-knowledged algorithm beats BIM, and BIM beats the random scheme.

Fig. 2 shows the evaluation result of BIM with varying T. We observe that all the three algorithms have the same informativeness when T is small (*e.g.*, less than 15), since the budget can afford all the users in each round. As T grows, full-knowledged algorithm achieves the best performance, since it is exploration-free and thus can utilize the total budget in the most efficient way. Different to the curves in Fig. 1, the total achieved informativeness keeps increasing and does not tend to saturate. This is because that with T increases, the submodularity of the informativeness function always allows us to reschedule some selected users to the new rounds without decreasing the total informativeness.

Fig. 3 shows the influence of the total budget in user selections. It varies the total budget B with fixed n and fixed T. We observe that the obtained informativeness first grows as B increases. That is because that more budget results in more selected users and thus more obtained sensing information. As the budget keeps increasing, the saturation will be reached when all the users are selected in every round.

We also consider scenarios where B varies with the increase of n or T. In Fig. 4, n ranges from 5 to 100, with B set to $50n$. We can see that the achieved informativeness increases as n grows. This is because that more budget allows more users to be selected. We also observe that the growth of informativeness tends to get slow as n increases, *e.g.*, 50 users can achieve over 300 informativeness by Full, while 100 users only achieve less than 500. This is due to the submodularity of the informativeness function. Besides, BIM achieves superior performance to the random selection scheme, and close performance to the full-knowledged algorithms.

Fig. 5 varies T from 5 to 100 and sets B to $200T$. We can see that the achieved informativeness is nearly proportional to T. This is supported by an intuitive observation that under the varying budget $200T$, if we keep selecting the best users in

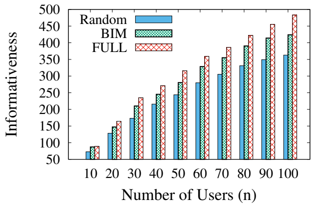

Fig. 4. Varying n ($B = 50n$)

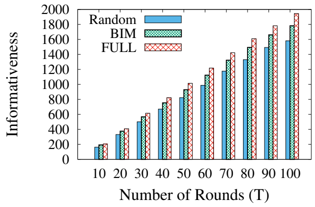

Fig. 5. Varying T ($B = 200T$)

the upcoming rounds, the achieved informativeness is linear to the round number. It can be seen that BIM achieves close informativeness to the full-knowledged algorithm.

VI. Related Work

Crowdsensing has been extensively studied in recent years. Yang *et al.* [5] studied a platform-centric model and a user-centric model in crowdsensing, and provided incentive mechanisms for them respectively. Zhao *et al.* [4] studied the budgeted online crowdsensing, and proposed two online budget feasible mechanisms. Jin *et al.* [23] incorporated a quality of information metric into the design of truthful combinatorial mechanisms for crowdsensing systems. Peng *et al.* [6] considered the quality-based pricing in crowdsensing, and presented an EM algorithm to quantify users' data qualities. However, to the best of our knowledge, none of the previous work in crowdsensing has addressed the problem of selecting most informative users, nor considered users' unknown costs.

The Gaussian process and mutual information-based model have been widely studied in sensor networks. In [2], the authors built a statistics model for sensor querying, based on the spatial and temporal correlations of environmental phenomena. Guestrin *et al.* [20] studied the most informative sensor placement problem, and adopted the mutual information-based sensor selection scheme. Krause *et al.* [24] further studied the problem of maximizing mutual information of sensors while minimizing the communication costs. However, the crowdsensing scenarios have clear differences from sensor networks, not only because users' costs are influenced by many factors and unknown a prior, but also due to the fact that autonomy of users makes us can only select users based on their current locations, instead of directly deploying users to sense some specific locations.

Multi-armed bandit problem was originally described by Robbins [9], after which many researchers have proposed different approaches (*e.g.*, [7], [10], [25]) to solve this problem. Chen *et al.* [11] studied a general framework for combinatorial multi-armed bandit problem. Budgeted MAB problems were preliminarily studied by [12], [14], [15]. They considered to optimize an additive objective function, where the costs are known and under a budget constraint. Different from the previous works, in this paper, the optimization function is submodualr, the number of selected users each round is uncertain, and users' sensing costs are unknown a priori.

VII. Conclusion

In this paper, we have considered the budgeted informativeness maximization problem in crowdsensing. We have proposed a mutual information-based criteria to quantify the informativeness. We have provided a multi-rounded approximation algorithm that achieves $(1 - 1/e)/2$ approximation ratio for full-knowledged scenarios. Without prior knowledge of users' costs, a MAB algorithm, namely BIM, has been proposed to efficiently select users. Our evaluation results have shown that our proposed algorithm achieves desirable performance in terms of the total obtained informativeness.

References

[1] K. L. Huang, S. S. Kanhere, and W. Hu, "On the need for a reputation system in mobile phone based sensing," *Ad Hoc Networks*, vol. 12, pp. 130–149, 2014.
[2] A. Deshpande, C. Guestrin, S. R. Madden, J. M. Hellerstein, and W. Hong, "Model-driven data acquisition in sensor networks," in *VLDB*, 2004.
[3] Y. Singer, "Budget feasible mechanisms," in *FOCS*, 2010.
[4] D. Zhao, X.-Y. Li, and H. Ma, "How to crowdsource tasks truthfully without sacrificing utility: Online incentive mechanisms with budget constraint," in *INFOCOM*, 2014.
[5] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing," in *MobiCom*, 2012.
[6] D. Peng, F. Wu, and G. Chen, "Pay as how well you do: A quality based incentive mechanism for crowdsensing," in *MobiHoc*, 2015.
[7] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2-3, pp. 235–256, 2002.
[8] T. M. Cover and J. A. Thomas, *Elements of information theory*. John Wiley & Sons, 2012.
[9] H. Robbins, "Some aspects of the sequential design of experiments," in *Herbert Robbins Selected Papers*. Springer, 1985, pp. 169–177.
[10] J. Vermorel and M. Mohri, "Multi-armed bandit algorithms and empirical evaluation," in *ECML*, 2005.
[11] W. Chen, Y. Wang, and Y. Yuan, "Combinatorial multi-armed bandit: General framework and applications," in *ICML*, 2013.
[12] L. Tran-Thanh, A. Chapman, J. E. Munoz De Cote Flores Luna, A. Rogers, and N. R. Jennings, "Epsilon–first policies for budget–limited multi-armed bandits," in *AAAI*, 2010.
[13] C.-J. Ho and J. W. Vaughan, "Online task assignment in crowdsourcing markets." in *AAAI*, 2012.
[14] L. Tran-Thanh, S. Stein, A. Rogers, and N. R. Jennings, "Efficient crowdsourcing of unknown experts using multi-armed bandits," in *ECAI*, 2012.
[15] L. Tran-Thanh, A. Chapman, A. Rogers, and N. R. Jennings, "Knapsack based optimal policies for budget-limited multi-armed bandits," *arXiv preprint arXiv:1204.1909*, 2012.
[16] W. Ding, T. Qin, X.-D. Zhang, and T.-Y. Liu, "Multi-armed bandit with budget constraint and variable costs," in *AAAI*, 2013.
[17] A. Biswas, S. Jain, D. Mandal, and Y. Narahari, "A truthful budget feasible multi-armed bandit mechanism for crowdsourcing time critical tasks," in *AAMAS*, 2015.
[18] A. Singla and A. Krause, "Truthful incentives in crowdsourcing tasks using regret minimization mechanisms," in *WWW*, 2013.
[19] S. Yang, F. Wu, S. Tang, L. Tie, X. Gao, L. Kong, and G. Chen. (2016, May) Selecting most informative contributors with unknown costs for budgeted crowdsensing. Technical Report. [Online]. Available: https://www.dropbox.com/s/rry8t2c8ti70crk/report-IWQoS.pdf?dl=0
[20] C. Guestrin, A. Krause, and A. P. Singh, "Near-optimal sensor placements in gaussian processes," in *ICML*, 2005.
[21] A. Krause and D. Golovin, "Submodular function maximization," *Tractability: Practical Approaches to Hard Problems*, vol. 3, p. 19, 2012.
[22] S. Khuller, A. Moss, and J. S. Naor, "The budgeted maximum coverage problem," *Information Processing Letters*, vol. 70, no. 1, pp. 39–45, 1999.
[23] H. Jin, L. Su, D. Chen, K. Nahrstedt, and J. Xu, "Quality of information aware incentive mechanisms for mobile crowd sensing systems," in *Proceedings of the 16th ACM International Symposium on Mobile Ad Hoc Networking and Computing*. ACM, 2015, pp. 167–176.
[24] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg, "Near-optimal sensor placements: Maximizing information while minimizing communication cost," in *IPSN*, 2006.
[25] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The nonstochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, pp. 48–77, 2002.