A Real-Time Framework for Task Assignment in Hyperlocal Spatial Crowdsourcing

LUAN TRAN and HIEN TO, University of Southern California

LIYUE FAN, University at Albany, SUNY

CYRUS SHAHABI, University of Southern California

Spatial Crowdsourcing (SC) is a novel platform that engages individuals in the act of collecting various types of spatial data. This method of data collection can significantly reduce cost and turnover time and is particularly useful in urban environmental sensing, where traditional means fail to provide fine-grained field data. In this study, we introduce hyperlocal spatial crowdsourcing, where all workers who are located within the spatiotemporal vicinity of a task are eligible to perform the task (e.g., reporting the precipitation level at their area and time). In this setting, there is often a *budget* constraint, either for every time period or for the entire campaign, on the number of workers to activate to perform tasks. The challenge is thus to maximize the number of assigned tasks under the budget constraint despite the dynamic arrivals of workers and tasks. We introduce a taxonomy of several problem variants, such as *budget-per-time-period* vs. *budget-per-campaign* and *binary-utility* vs. *distance-based-utility*. We study the hardness of the task assignment problem in the *offline* setting and propose *online* heuristics which exploit the spatial and temporal knowledge acquired over time. Our experiments are conducted with spatial crowdsourcing workloads generated by the SCAWG tool, and extensive results show the effectiveness and efficiency of our proposed solutions.

CCS Concepts: • **Information systems** → **Crowdsourcing**; **Geographic information systems**; • **Human-centered computing** → **Ubiquitous and mobile computing**;

Additional Key Words and Phrases: Spatial crowdsourcing, crowdsensing, participatory sensing, GIS, online task assignment, budget constraints

ACM Reference format:

Luan Tran, Hien To, Liyue Fan, and Cyrus Shahabi. 2017. A Real-Time Framework for Task Assignment in Hyperlocal Spatial Crowdsourcing. *ACM Trans. Intell. Syst. Technol.* 9, 3, Article 37 (January 2018), 26 pages. <https://doi.org/10.1145/3078853>

1 INTRODUCTION

With the ubiquity of smart phones and improvements in wireless network bandwidth, every person with a mobile phone can now act as a multimodal sensor, collecting and sharing various types

L. Tran and H. To contributed equally this work.

This research has been funded by NSF grants IIS-1320149, CNS-1461963, the USC Integrated Media Systems Center, and the University at Albany. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any of the sponsors such as NSF.

Authors' addresses: L. Tran, 1018 1/2 W 23rd St, Los Angeles, CA 90007; email: luantran@usc.edu; H. To, 1711 Summit Ave APT 106, Seattle, WA 98122; email: hiento@amazon.com; L. Fan, 1400 Washington Avenue, Albany, NY 12222; email: liyuefan@albany.edu; C. Shahabi, University of Southern California, Computer Science Department, SAL 300 Los Angeles, CA 90089-0781; email: shahabi@usc.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

© 2017 ACM 2157-6904/2017/01-ART37 \$15.00

<https://doi.org/10.1145/3078853>

Fig. 1. Hyperlocal spatial crowdsourcing framework.

of high-fidelity spatiotemporal data instantaneously. In particular, crowdsourcing for weather information has become popular. With a few recent apps, such as mPING¹ and WeatherSignal², individual users can report weather conditions, air pollution, noise levels, and more. In fact, Dorminey (2014) regards crowdsourcing as “the future of weather forecasting.”

Through our collaboration with the Center for Hydrometeorology and Remote Sensing (CHRS)³ at the University of California, Irvine, we have developed a mobile app, iRain⁴ (iRa 2016), to perform *spatial crowdsourcing* for precipitation information. Unlike other weather crowdsourcing apps, iRain allows CHRS researchers to *request* rainfall information at specific locations and times where their global satellite precipitation estimation technologies⁵ fail to provide real-time, fine-grained data. Individual iRain users around those locations can *respond* to those requests by reporting rainfall observations (e.g., heavy/medium/light/none).

In general, spatial crowdsourcing (SC) (Kazemi and Shahabi 2012) offers an effective data collection platform where data requesters can create spatial tasks dynamically and workers are assigned to tasks based on their locations. Figure 1 depicts the architecture of iRain. A requester issues a set of rainfall observation tasks to the SC-server (Step 1) where each task corresponds to a specific geographical extent (e.g., a circle). The workers continuously update their locations to the SC-server when they become available for performing tasks (Step 0). Subsequently, the SC-server crowdsources the tasks among the workers in the task regions and sends the collected data back to the requester (Steps 2, 3).

One major difference from existing SC paradigms (Kazemi and Shahabi 2012; He et al. 2014; To et al. 2015; Xiao et al. 2015; Guo et al. 2016) is that workers in our paradigm do not need to travel to the exact task locations (e.g., to the centers of the circular regions) and are eligible to perform tasks as long as they are in close spatiotemporal vicinity of the tasks (e.g., enclosed in the circular regions).⁶ We denote this new paradigm *Hyperlocal Spatial Crowdsourcing*. The reason is twofold. Without requiring workers to travel physically, our paradigm lowers the threshold for worker participation and will potentially yield faster response. Furthermore, the requested data (e.g., rainfall or temperature) exhibit spatiotemporal continuity in measurement. Therefore,

¹<http://mping.nssl.noaa.gov/>.

²<http://weathersignal.com>.

³<http://chrs.web.uci.edu/>.

⁴<https://itunes.apple.com/us/app/irain-uci/id982858283>.

⁵<http://hydis.eng.uci.edu/gwadi/>.

⁶Tasks that require workers to physically travel to task locations (e.g., taking a picture of an event) are not considered in our problem setting.

observations obtained at nearby locations (e.g., within a certain distance to the task location and close to the requested time) are sufficient to fulfil the task. For example, workers *B* and *C* in Figure 1 are both eligible to report precipitation level at the University of Southern California (USC), and worker *A*, who becomes available 5 minutes later, is also qualified. The acceptable ranges of space and time can be specified by data requesters, from which the SC-server can find the set of eligible workers for each task.

The SC-server operates to maximize fulfilled tasks for revenue. Therefore, it cannot assign tasks to an unlimited number of workers due to budget considerations. In Hyperlocal SC, the *budget* represents the payment to each selected worker upon task completion, or the communication cost for sending/receiving task notifications between the SC-server and each selected worker. Furthermore, it is not necessary to select many workers for overlapping tasks. For example, in Figure 1, the observation of worker *A* can be used for precipitation tasks at both USC and Los Angeles downtown (shown in two circles).

The goal of our study is to maximize the number of assigned tasks on the SC-server where only a given number of workers can be selected over a time period or during the entire campaign (i.e., under “budget” constraints). When tasks and workers are known *a priori*, we can reduce the task assignment problem to the classic *Maximum Coverage Problem* and its variants. However, the main challenge with SC comes from the dynamism of the arriving tasks and workers, which renders an optimal solution infeasible in the online scenario. In Figure 1, the SC-server is likely to activate worker *D* and either worker *B* or *C* for the two tasks, respectively, without knowing that a more favorable worker *A* is qualified for both tasks and will arrive in the near future. Previous heuristics in the literature (Kazemi and Shahabi 2012; To et al. 2015; Deng et al. 2016; Guo et al. 2016) do not consider the vicinity of tasks in space and time or the budget, and thus cannot be applied to Hyperlocal SC.

The contributions of this article are as follows⁷ (i) We provide a formal definition of Hyperlocal SC where the goal is to maximize task coverage under budget constraints. We introduce a taxonomy to classify several problem variants (e.g., given a budget constraint for each time period [*fMTC*] vs. for the entire campaign [*dMTC*]). We show that both *fMTC* and *dMTC* variants are NP-hard in the offline setting. (ii) In the online setting, we propose several heuristics for real-time task assignment. When a budget constraint is given for each time period (*fMTC*), local heuristics (*Basic*, *Temporal*, and *Spatial*) are developed to select workers within each time period. When a budget is given for the entire campaign (*dMTC*), we devise an adaptive strategy based on the contextual bandit to dynamically allocate the total budget to a number of time periods. (iii) When the utility of an assigned task is considered, we introduce two distance-based utility models to measure the assignment quality based on worker-task distance, which can be integrated with any previously developed heuristics. To avoid overloading workers, we introduce a multiobjective variant to minimize the repetitive activation of the same worker. Online solutions based on a genetic algorithm and adaptive budget allocation are developed for *fMTC* and *dMTC* scenarios, respectively. (iv) We conduct extensive experiments with workbench datasets generated from real-world location check-ins. The empirical results confirm that our heuristics are efficient and effective in assigning hyperlocal tasks in a real-time manner.

The remainder of this article is organized as follows. Section 2 reviews the related work. Section 3 provides notations and a taxonomy for the Hyperlocal SC problem. We prove in Section 4 that offline task assignment in Hyperlocal SC with budget constraints is NP-hard. In Section 5, we study two problem variants in the online setting. Section 6 discusses the multiobjective optimization variant to mitigate worker overloading, and Section 7 describes the integration of distance-based

⁷This article is an extension of a short paper that in To et al. (2016b).

task utility models. We report our experimental results in Section 8, provide discussion in Section 9, and conclude the article in Section 10.

2 RELATED WORK

Spatial crowdsourcing can be deemed one of the main enablers of urban computing's applications such as monitoring traffic information and air pollution (Zheng et al. 2014; Ji et al. 2016). Only recently has SC gained popularity in both the research community and industry (e.g., TaskRabbit, Gigwalk). A recent study (To et al. 2015) distinguishes SC from related fields, including generic crowdsourcing, participatory sensing, volunteered geographic information, and online matching. Research efforts in SC have focused on different aspects, such as *task assignment* (e.g., Kazemi and Shahabi (2012), Tong et al. (2016), Liu et al. (2016b), Cheng et al. (2016), and Hu et al. (2016)), *task scheduling* (e.g., Deng et al. (2016), Li et al. (2015), and Sales Fonteles et al. (2016)), *quality control* and *trust* (e.g., Kazemi et al. (2013) and Cheng et al. (2015)), *privacy* (e.g., To et al. (2014, 2017) and Wang et al. (2017)), and *incentive mechanism* (e.g., Gao et al. (2015), Kandappu et al. (2016), and Zhang et al. (2017)). The authors in Kazemi and Shahabi (2012) proposed a task assignment problem whose goal is to maximize the number of assigned tasks. Requesters may want to crowdsource a spatial *complex* task that requires multiple workers at different locations to collectively perform several spatial subtasks (Dang et al. 2013; Zhang et al. 2016; Gao et al. 2017). In To et al. (2014), the authors introduce the problem of protecting worker location privacy in SC. A framework is proposed to ensure differentially private protection guarantees without significantly affecting the effectiveness and efficiency of the SC system. Another study proposes a differentially private incentive mechanism in a mobile crowd sensing system (Jin et al. 2016). In such systems where workers bid tasks, a worker's bid may reveal her interests and knowledge base. Thus, the proposed method preserves the privacy of each worker's bid against other honest-but-curious workers. However, unlike our study, this work does not focus on challenges that are unique to SC, such as spatial task allocation. Meanwhile, the authors in Pournajaf et al. (2014) propose a two-phase framework whose objective is to match a set of spatial tasks to a set of workers given the workers' cloaking regions such that task assignment is maximized while satisfying the travel budget constraint of each worker. The trust issues in SC have been studied in Kazemi et al. (2013), where one solution is having tasks performed redundantly by multiple workers. Recently, in Cheng et al. (2016), the reliability of task assignment is measured in term of both the confidence of task completion and the diversity quality of the tasks. However, the trust and reliability issues of workers are beyond the scope of our work; if there are multiple reports for one task, the SC-server will simply send all available reports to the task requester. It is worth noting that we assume the workers would respond to their assigned tasks, which is a common assumption in the server-assigned mode of SC (Kazemi and Shahabi 2012). There have been studies that relax this assumption by associating to each worker a probability to perform an assigned task (e.g., Cheng et al. (2016) and To et al. (2014)). However, this consideration is not the focus of our article. In addition, the time required for the workers responding to an assigned task is negligible (e.g., a few seconds with the iRain application) when compared to the deadline of each task (e.g., one day).

Online Spatial Task Assignment: There have been extensive studies regarding task assignment in generic crowdsourcing (e.g., Venanzi et al. (2013) and Tran-Thanh et al. (2013)). However, unlike our study that focuses on the spatiotemporal aspects of task assignment, these studies focus on task assignment in general crowdsourcing rather than SC. Recent studies that are closely related to ours include Tong et al. (2016) and ul Hassan and Curry (2014). Both of them study the online spatial task assignment problem; however, they differ from our work in terms of the problem setting and objectives. First, in our problem, the report of a worker can be used for multiple tasks as long as their geographical extents cover the worker's location. Thus, our focus is

worker selection rather than matching of workers to tasks as in Tong et al. (2016) and ul Hassan and Curry (2014). Second, the objectives in these studies are, respectively, to maximize the number of assigned tasks and to maximize the total utility score of the worker–task matches; whereas our framework considers both kinds of utility, each is jointly combined with other real-world considerations, such as leveraging historical workload and minimizing worker overloading. In addition, our study maximizes the utility of assignment under budget constraints while others focus solely on maximizing the utility.

Budgeted Spatial Task Assignment: There have been recent studies on matching workers with tasks under budget constraints. In Tran-Thanh et al. (2013), the authors propose Crowd-Budget, an agent-based budget allocation algorithm that divides a given budget among different tasks in order to achieve low estimation error (of the estimated answers for a set of tasks). This study does not consider the challenges of spatial task assignment, where workers and tasks can come and go at any time, and users are not aware of their locations until their arrival time. The study in Miao et al. (2016) also differs from ours. In our problem, workers do not need to travel and report sensed value at their current locations. In contrast, in Miao et al. (2016), workers do need to travel to the task locations, which may take a long time in rush hour. Consequently, the workers may reject their assigned tasks. It is worth noting that these studies focus on the worker–task matching problem; however, our aim is to select the best workers to maximize task coverage. Similar to Zhang et al. (2015) and Tran-Thanh et al. (2013), the budget in Miao et al. (2016) refers to payment to workers, while in our study, we consider budget as the number of workers to select and focus on allocating a total budget across multiple time periods.

Worker Selection: Several works studied the problem of selecting workers with budget constraints (Song et al. 2014; Zhang et al. 2014). However, those studies focus on the offline participant selection problem while our focus is to propose online solutions. Furthermore, the problem settings in those studies differ from ours in several aspects. Sensing tasks in Song et al. (2014) are represented by non-overlapping regions, while tasks in our study can overlap spatially and are thus more challenging for optimization. The authors in Zhang et al. (2014) studied the problem of selecting a minimum number of workers to minimize the overall incentive payment while satisfying a probabilistic coverage requirement; however, in our problem, the number of workers to be selected is constrained by a predefined budget. Our work is also closely related to the problem of matching workers with tasks (He et al. 2014; Xiao et al. 2015). Particularly, in He et al. (2014), the authors studied the problem of task allocation that maximizes the reward of the SC platform given a time constraint for each worker. Recently, in Xiao et al. (2015), a task assignment problem that minimizes the average makespan of all assigned tasks was proposed. Unlike these studies, SC workers in our setting need not travel to task locations. Furthermore, our aim is different from the aforementioned studies, which is to maximize task coverage.

3 PRELIMINARIES

We first introduce concepts and notations used in this article. A **task** is a query of certain hyperlocal information, such as precipitation level at a particular location and time. For simplicity, we assume that the result of a task is in the form of a numerical value (e.g., $0 = \text{rain}, 1 = \text{snow}, 2 = \text{none}$).⁸ Specifically, every task comes with a predefined region where any enclosed worker can report data for that task. In this article, we define each task region as a circular space centered at the task location; however, task region can be extended to other shapes such as a polygon or used to represent geography such as a district, city, county, and the like. Moreover, each task also specifies a valid *time interval* during which users can provide data.

⁸Remote sensing techniques based on satellite images cannot differentiate between rain and snow.

Fig. 2. Graphical example of worker-task coverage ($\delta = 2$). Subscripts represent time periods while superscripts mean ids.

Definition 3.1 (Task). A task t of form $\langle l, r, s, \delta \rangle$ is a query at location l , which can be answered by workers within a circular space centered at l with radius r . The parameter δ indicates the duration of the query: It is requested at time s and can be answered until time $s + \delta$.

We refer to $s + \delta$ as the “deadline” of task t . A task expires if it has not been answered before its deadline. Figure 2(a) shows the regions of six tasks, $t_1^1, t_1^2, \dots, t_1^6$. All tasks expire at time period 2 (i.e., they can be deferred to time period 2), represented by the dashed circles in Figure 2(b).

A **worker** can accept task assignments when he is *online*.

Definition 3.2 (Worker). A worker w of form $\langle id, l \rangle$, is a carrier of a mobile device who can accept spatial task assignments. The worker can be uniquely identified by his *id* and his location is at l .

Intuitively, a worker is eligible to perform a task if his location is enclosed in the task region. In Figure 2(a), w_1^1 is eligible to perform t_1^1, t_1^2 , and t_1^3 , while w_2^1 is qualified for t_1^4, t_1^5 and t_1^6 . Furthermore, a worker’s report to one task can also be used for all other unexpired tasks whose task regions enclose the worker. As in Figure 2(b), w_2^1 is eligible to perform t_1^1 and t_1^5 , which are deferred from time 1.

Let $W_i = \{w_i^1, w_i^2, \dots\}$ denote the set of available workers at time s_i and $T_i = \{t_i^1, t_i^2, \dots\}$ denote the set of available tasks, including tasks issued at time s_i and previously issued unexpired tasks. Next, we define the notions of **worker-task coverage** and **coverage instance sets**.

Definition 3.3 (Worker-Task Coverage). Given $w_i^j \in W_i$, let $C(w_i^j) \subset T_i$ denote the task coverage set of w_i^j , such that for every $t_i^k \in C(w_i^j)$,

$$s_i < t_i^k.(s + \delta) \quad (1)$$

$$\|w_i^j.l - t_i^k.l\|_2 \leq t_i^k.r \quad (2)$$

We also say the worker w_i^j covers the tasks $t_i^k \in C(w_i^j)$. An example of coverage in Figure 2(a) is $C(w_1^1) = \{t_1^1, t_1^2, t_1^3\}$.

Definition 3.4 (Coverage Instance Set). At time s_i , the coverage instance set, denoted by I_i is the set of worker-task coverage of form $\langle w_i^j, C(w_i^j) \rangle$ for all workers $w_i^j \in W_i$.

The coverage instance sets for the example in Figure 2 are illustrated in Table 1. For simplicity, we now assume the utility of a specific task assignment is *binary* within the task region and before

Table 1. The Coverage Instance Set of the Example in Figure 2

Time	Coverage Instance Sets
1	$\{(w_1^1, \langle t_1^1, t_1^2, t_1^3 \rangle), (w_1^2, \langle t_1^1, t_1^4, t_1^5, t_1^6 \rangle)\}$
2	$\{(w_2^1, \langle t_1^5, t_1^6 \rangle)\}$

the deadline. That is, assignment to any worker within a task region before the deadline has utility 1 (1 successful assignment) and 0 otherwise. As a result, tasks t_1^5 and t_1^6 being answered by worker w_1^2 at time 1 is equivalent to them being answered by w_2^1 at time 2.

Again, the goal of our study is to maximize task assignment given a budget, despite the dynamic arrivals of tasks and workers. Now, we formally define the notion of a budget.

Definition 3.5 (Budget). Budget K is the maximum number of workers to select in a coverage instance set.

In practice, budget K can capture the *communication cost* the SC-server incurs to push notifications to selected workers (Step 3 in Figure 1) or the *rewards* paid to the workers.

With these, we formally define the hyperlocal crowdsourcing problem with budget constraint as follows:

Definition 3.6 (Problem). Given a set of workers W , a set of available tasks T , a budget constraint K , and a utility function $U(\cdot) \in \mathbb{R}$, find a subset of workers W' of W within the budget constraint such that the total utility of the covered tasks is maximized.

3.1 Problem Taxonomy

3.1.1 Budget-per-Time-Period vs. Budget-per-Campaign. In certain scenarios, the task requester may specify a budget constraint (i.e., the maximum number of workers to select) for **each time period** in a campaign (e.g., a day or a week). Given a set of time periods $\phi = \{s_1, s_2, \dots, s_Q\}$, a budget constraint K_i is specified for each s_i . The challenge is to decide which workers to select within each time period. On the other hand, the task requester could specify a budget constraint for the **entire campaign**. Given a set of time periods $\phi = \{s_1, s_2, \dots, s_Q\}$ and assuming L_i workers are selected for s_i , a budget constraint K is specified for the sum of L_i 's. The new challenge of this problem variant is to allocate the total budget K wisely over Q time periods.

The **choice** of the constraint model depends on the financial flexibility of the task requester. Furthermore, the budget-per-time-period model is a special case of the budget-per-campaign model. As a result, the utility of the budget-per-campaign solution is no worse than that of the budget-per-period solution for any problem instance.

3.1.2 Binary-Utility vs. Distance-Based-Utility. Considering the utility of assigned tasks, our problem can be classified into binary-utility and distance-based-utility variants. In the **binary-utility** model, a task can be assigned to any worker located within the task radius to achieve utility 1. Unassigned tasks will yield 0 utility. Therefore, the optimization objective is to maximize the total number of assigned tasks. However, for some applications, a worker who is closer to the task location may be “preferred” over other workers farther away. For example, in weather crowdsourcing applications (e.g., iRain (iRa 2016)), a closer worker can report more accurate rainfall data. The **distance-based-utility** model thus evaluates a task assignment to a specific worker with various distance functions.

3.1.3 Single-Objective vs. Multiobjective. In the **single-objective** problem formulation, we aim to maximize the total utility of assigned tasks. On the other hand, crowdsourcing applications may have more than one, sometimes conflicting, objectives to ensure long-term prosperity. For example, worker overloading can be a critical concern of the novel crowdsourcing platforms, in which only a few workers are frequently selected to optimize task assignments. Therefore, a **multiobjective** formulation can introduce a second objective to minimize the worker overloading phenomenon. The challenge is thus to find solutions considering the tradeoff between the two objectives.

3.1.4 Offline vs. Online. Orthogonal to the preceding dimensions, our problem can be further classified into offline and online variants. The **offline** variant selects workers with complete knowledge of task/worker arrivals during the entire campaign. Although this is not practical, studying the offline variant allows us to eliminate the hardness arising from the randomness of the online problem and focus on the optimization in a deterministic setting. In the **online** variant, assignments have to be made in real-time for the currently arriving tasks/workers without complete knowledge of future arrivals. While it is more fitting for crowdsourcing applications, it is also intuitively more challenging: It is uncertain in nature when and where future tasks and workers may appear. Thus, effective worker selection must optimize the objective(s) in the long run.

The majority of this article will focus on the online, binary-utility, single-objective problem with both per-time-period and per-campaign budget constraints. We will also show how to extend our solutions to the distance-based utility model as well as the multiobjective problem.

4 HARDNESS OF THE PROBLEM

In this section, we study the problem complexity of task assignment with budget constraint in hyperlocal SC. We show that two offline variants (i.e., budget-per-time-period vs. budget-per-campaign) of the problem are NP-hard and propose online heuristics in the next section.

4.1 Fixed Budget *fMTC*

PROBLEM 1 (FIXED-BUDGET MAXIMUM TASK COVERAGE). *Given a set of time periods $\phi = \{s_1, s_2, \dots, s_Q\}$ and a budget K_i for each s_i , the fixed-budget maximum task coverage (*fMTC*) problem is to select a set of workers L_i at every s_i , such that the total number of covered tasks $|\bigcup_{i=1}^Q \bigcup_{w_i^j \in L_i} C(w_i^j)|$ is maximized and $|L_i| \leq K_i$.*

This optimization problem is challenging since each worker is eligible for a subset of tasks. The fact that a task can be deferred to future time periods further adds to the complexity of the problem. With the following theorem, we prove that *fMTC* is NP-hard by a reduction from the *maximum coverage with group budgets* (MCG) constraints problem (Chekuri and Kumar 2004). MCG is motivated by the maximum coverage problem (MCP) (Feige 1998). Consider a given I_g , we are given the subsets $S = \{S_1, S_2, \dots, S_m\}$ of a ground set X and the disjoint sets $\{G_1, G_2, \dots, G_l\}$. Each G_i , namely a *group*, is a subset of $S = \{S_1, S_2, \dots, S_m\}$. With MCG, we are given an integer k and an integer bound k_i for each group G_i . A solution to I_g is a subset $H \subset S$ such that $|H| \leq k$ and $|H \cap G_i| \leq k_i$ for $1 \leq i \leq l$. The objective is to find a solution such that the number of elements of X covered by the sets in H is maximized. MCP is the special case of MCG. Since MCP is known to be strongly NP-hard (Feige 1998), by restriction, MCG is also NP-hard.

THEOREM 1. *fMTC is NP-hard.*

PROOF. We prove the theorem by a reduction from MCG (Chekuri and Kumar 2004). That is, given an instance of the MCG problem, denoted by I_g , there exists an instance of the *MTC* problem,⁹

⁹In this section, *MTC* refers to fixed-budget *MTC* for short.

denoted by I_t , such that the solution to I_t can be converted to the solution of I_g in polynomial time. The reduction has two phases, *transforming* all workers/tasks across the entire campaign to a bipartite graph, and *mapping* from MCG to MTC. First, we layout the tasks and workers as two sets of vertices in a bipartite graph in Figure 2(c). A worker w_i^j can cover a task t_i^k if both spatial and temporal constraints hold; see Equations (2) and (1), respectively. In Figure 2(c), w_2^1 can cover t_1^4 and t_2^5 , which are deferred from s_1 to s_2 , represented by the dashed line.

Thereafter, MTC can be stated as follows. Selecting the maximum K_i workers per group, each group represents a time period, such that the number of covered tasks is maximized (i.e., $|L_i| \leq K_i$). To reduce I_g to I_t , we show a mapping from I_g components to I_t components. For every element in the ground set X in I_g , we create a task t_i^j ($1 \leq j \leq |X|$). Also, for every set in S , we create a worker w_i^j with $C(w_i^j) = S_j$ ($1 \leq j \leq m$). Consequently, to solve I_t , we need to find a subset $L_i \subset W_i$ workers of maximum size K_i in each group whose coverage is maximized. Clearly, if an answer to I_t is the set L_i ($1 \leq i \leq Q$), the answer to I_g will be the set $H \subset S$ of maximum coverage such that $|H| \leq k = \sum_{i=1}^Q K_i$ and $|H \cap G_i| \leq k_i = K_i$ for $1 \leq i \leq Q$.

As the transformation is bounded by the polynomial time to construct the bipartite graph, this completes the proof. $\square$

By a reduction from the MCG problem, we can now use any algorithm that computes MCG to solve the MTC problem. The greedy algorithm in Chekuri and Kumar (2004) provides 0.5-approximation for MCG. For example, the greedy solution in Figure 2(c) is $\{w_1^1, w_2^1\}$. However, the approximation ratio only holds in the offline scenario where the server knows *a priori* the coverage instance set for every time period.

4.2 Dynamic Budget dMTC

PROBLEM 2 (DYNAMIC-BUDGET MAXIMUM TASK COVERAGE). *The dynamic-budget maximum task coverage problem (dMTC), is similar to fMTC, except the total budget K is specified for the entire campaign (i.e., $\sum_{i=1}^Q |L_i| \leq K$).*

In the offline setting where the server is clairvoyant about future workers and tasks, we prove the dMTC problem is NP-hard by reduction from the MCP.

THEOREM 2. *dMTC is NP-hard.*

PROOF. We prove the theorem by a reduction from MCP. That is, given an instance of the MCP problem, denoted by I_m , there exists an instance of the MTC problem,¹⁰ denoted by I_t , such that the solution to I_t can be converted to the solution of I_m in polynomial time. The reduction includes two steps, *transforming* all workers/tasks across the entire campaign to a bipartite graph, and *mapping* from MCP to MTC. The first step is similar to that of Theorem 1, in which the workers and tasks from the entire campaign are transformed into a bipartite graph, as illustrated in Figure 2(c). The *mapping* step can be considered as a special case of the proof in Theorem 1, in which there exists only one group of all budget. As the transformation is bounded by the polynomial time to construct the bipartite graph and MCP is strongly NP-hard, this completes the proof. $\square$

The results of these solutions to the offline scenarios will be used as the upper bounds of the results to the online solutions to be discussed in Section 5.

¹⁰In this section, MTC refers to dynamic-budget MTC for short.

5 ONLINE TASK ASSIGNMENT

In this section, we focus on *online* variants: online *fMTC* when a budget constraint is given for each time period and online *dMTC* when the budget constraint is given for the entire campaign. We will introduce heuristics for each variant as follows.

5.1 Fixed Budget *fMTC*

In the online scenario where workers and tasks arrive dynamically, it becomes more challenging to achieve the global optimal solution for Problem 1. Since the server does not have prior knowledge about future workers and tasks, it tries to optimize task assignment locally at every time period. However, the optimization within every time period, similar to the MCP, is also NP-hard. A greedy algorithm (Feige 1998) was proposed to achieve an approximation ratio of 0.63 by choosing a set which contains the largest number of *uncovered* elements at each stage. This study shows that the greedy algorithm is the best possible polynomial time approximation algorithm for MCP. Next, we propose several greedy heuristics to solve the online *fMTC* problem; namely *Basic*, *Spatial*, and *Temporal*.

5.1.1 Basic Heuristic. The *Basic* heuristic solves the online *fMTC* problem by using the greedy algorithm (Hochbaum 1996) for every time period. At each stage, *Basic* selects the worker who covers the maximum number of *uncovered* tasks, depicted in Line 10 of Algorithm 1. For instance, in Figure 2(a), w_1^2 is selected at the first stage. At the beginning of each time period, Line 4 removes expired tasks from the previous time period. Line 5 adds unassigned, unexpired tasks to the current task set. Line 12 outputs the covered tasks C_i per time period, which will be used as the main performance metric in Section 8. The algorithm terminates when either running out of budget or when all the tasks are covered (Line 9).

ALGORITHM 1: BASIC ALGORITHM

```

1: Input: worker set  $W_i$ , task set  $T_i$ , budgets  $K_i$ 
2: Output: selected workers  $L_i$ 
3: For each time period  $s_i$ 
4:   Remove expired tasks  $U'_{i-1} \leftarrow U_{i-1}$ 
5:   Update task set  $T_i \leftarrow T_i \cup U'_{i-1}$ 
6:   Remove tasks that do not enclose any worker  $T'_i \leftarrow T_i$ 
7:   Construct worker set  $W_i$ , each  $w_i^j$  contains  $C(w_i^j)$ 
8:   Init  $L_i = \{\}$ , uncovered tasks  $R = T'_i$ 
9:   While  $|L_i| < K_i$  and  $|R| > 0$ 
10:    Select  $w_i^j \in W_i - L_i$  that maximize  $|C(w_i^j) \cap R|$ 
11:     $R \leftarrow R - C(w_i^j)$ ;  $L_i \leftarrow L_i \cup \{w_i^j\}$ 
12:   $C_i \leftarrow \bigcup_{w_i^j \in L_i} C(w_i^j)$ 
13:  Keep uncovered tasks  $U_i \leftarrow T'_i - C_i$ 

```

Basic can achieve fast task assignment by simply counting the number of tasks covered by each worker (Line 10). However, it treats all tasks equally, without considering the spatial and temporal information of each task (i.e., location and deadline). For example, a task located in a “worker-sparse” area may not be assigned in the future due to lack of nearby workers and thus should be assigned with higher priority at the current iteration. Similarly, tasks that are expiring soon should be assigned with higher priorities. Consequently, the priority of a worker is high if he covers a larger number of high-priority tasks. Next, we introduce two assignment heuristics that explicitly model the task priority given its spatial and temporal characteristics.

Fig. 3. Approximation of Task Entropy.

5.1.2 Temporal Heuristic. One approach to prioritizing tasks is by considering their temporal urgency. The intuition is that a task which is further away from its deadline is more likely to be covered in the future, and vice versa. As a result, near-deadline tasks should have higher priorities to be assigned than others. Consequently, a worker who covers a large number of soon-to-expire tasks should be preferred for selection. Based on this intuition, we model the priority of a worker w_i^j based on the remaining time of each task he covers as follows:

$$\text{priority}(w_i^j) = \sum_{t_i^k \in C(w_i^j) \cap R} \frac{1}{t_i^k \cdot (s + \delta) - i}. \quad (3)$$

The *Temporal* heuristic adapts *Basic* by selecting the worker with maximum *priority* at each stage. For instance, consider two workers w_1^1 and w_1^2 at time s_1 , where $C(w_1^1) = \{t_1^1, t_1^2\}$ and $C(w_1^2) = \{t_1^3\}$. Suppose both t_1^1 and t_1^2 expire in 5 time periods and t_1^3 expires in 2 time periods. The *Temporal* heuristic chooses w_1^2 over w_1^1 because their priorities are 0.5 and 0.4, respectively. To implement *Temporal*, Line 10 in Algorithm 1 can be updated to select the worker with maximum priority, as defined in Equation (3). We will empirically evaluate all heuristics in Section 8.

5.1.3 Spatial Heuristic. To maximize task assignment in the long term, we also consider the “popularity” of a task location as an indicator of whether the task can be assigned to future workers. Accordingly, we can spend the budget for the current time period to assign those tasks which can be only covered by existing workers. The “popularity” of a task region can be measured using Location Entropy (Cranshaw et al. 2010), which captures the diversity of visits to that region. A region has high entropy if many workers visit with equal probabilities. In contrast, a region has low entropy if there are only a few workers visiting. We define the *entropy* of a given task as follows.

For task t , let O_t be the set of visits to the task region $R(t.l, r)$. Let W_t be the set of distinct workers who visited $R(t.l, r)$, and let $O_{w,t}$ be the set of visits that worker w made to $R(t.l, r)$. The probability that a random draw from O_t belongs to $O_{w,t}$ is $P_t(w) = \frac{|O_{w,t}|}{|O_t|}$. The entropy of t is computed as follows:

$$RE(t) = - \sum_{w \in W_t} P_t(w) \times \log P_t(w). \quad (4)$$

For efficient evaluation, $RE(t)$ can be approximated by aggregating the entropies of 2D grid cells within the task region $R(t.l, r)$, and the cell entropies can be precomputed using historical data. Since any worker located inside $R(t.l, r)$ can perform task t , t is likely to be covered in the future as long as one grid cell inside $R(t.l, r)$ is “popular” among workers. Figure 3 illustrates the precomputation of the entropy of task t . When a task arrives, we first identify the grid cell that

encloses the task location (i.e., the white cell in the center) and slightly adjust the task region (solid circle) to be centered at the white cell (dashed circle). We approximate the task entropy by the entropy of the dashed circle, which can be computed. This is because the dashed circle is solely determined by the white cell and radius r . To further speed up the precomputation of all possible combinations of the cell and the radius, we approximate the dashed circle by a set of cells whose centers are within the circle. With the entropy of every task covered by worker w_i^j , his priority can be calculated as follows:

$$\text{priority}(w_i^j) = \sum_{t_i^k \in C(w_i^j) \cap R} \frac{1}{1 + RE(t_i^k)}. \quad (5)$$

Note that the constant 1 is needed to avoid division by zero. Consequently, the *Spatial* heuristic greedily selects the worker with maximum *priority* at each stage. Line 10 in Algorithm 1 can be modified to reflect the spatial priority of each worker.

5.2 Dynamic Budget *dMTC*

The second problem variant we study is more general, where a budget constraint is given for the entire campaign. This relaxation often results in higher task coverage. For example, in Figure 2, if budget 1 is given at every time period, we select w_1^1 and w_2^1 and obtain the coverage of 5. However, the dynamic-budget variant yields higher coverage of 6 by selecting w_1^1 and w_1^2 at time 1. Next, we study the problem complexity in the offline scenario and propose adaptive budget allocation strategies for the online scenario.

The challenge of the online *dMTC* problem is to allocate the overall budget K over Q time periods ($K \geq Q$) optimally, despite the dynamic arrivals of workers and tasks. Next, we introduce several budget allocation strategies. Once a budget is allocated to a particular time period, we can adopt previously proposed heuristics (*Basic*, *Spatial*, *Temporal*) to select the best worker.

The simplest strategy, namely *Equal*, equally divides K to Q time periods; each time period has K/Q budget, and the last time period obtains the remainder. However, *Equal* may overallocate budget to those time periods with small numbers of tasks. Another strategy is to allocate a budget to each time period proportional to the number of available tasks at that time period; that is, $\frac{|T_i|}{|T|}K$, where $|T|$ is the total number of tasks. However, $|T|$ is not known *a priori*. Furthermore, we may still overallocate budget to any time period with large $|T_i|$ if none of the tasks can be covered by any workers (or all the tasks can be covered by 1 worker). We cannot allocate budget optimally without looking at the coverage instance set at each time period.

5.2.1 Adaptive Budget Allocation. To maximize task assignment, we need to adaptively allocate the overall budget and consider the “return” of selecting every worker (i.e., the worker priority), given the dynamic coverage instance set at every time period. We define the following two notions. **Delta budget**, denoted as δ_K , captures the current status of budget utilization compared to a baseline budget strategy $\{K^{base}[t], t = 1, \dots, Q\}$ (e.g., the *Equal* strategy). Given a certain baseline $\{K^{base}[t]\}$, δ_K is the difference between the cumulative baseline budget and the actual budget spent up to time period s_i . Formally, at any time period s_i ,

$$\delta_K = \sum_{t=1}^i (K^{base}[t]) - K_{used}. \quad (6)$$

A positive δ_K indicates budget is underutilized, and vice versa. Another notion is **delta gain**, denoted as δ_λ , which represents the return of a worker currently being considered (λ_l) compared

Fig. 4. Adaptive budget allocation based on contextual ϵ -greedy.

to the ones selected in the past ($\overline{\lambda_{l-1}}$). Formally,

$$\delta_\lambda = \lambda_l - \overline{\lambda_{l-1}}, \quad (7)$$

where λ_l is the gain of the current worker, calculated by any previously proposed local heuristic (as $|priority(w_i^j)|$). $\overline{\lambda_{l-1}}$ is the average gain of previously added workers (i.e., $\overline{\lambda_{l-1}} = \frac{1}{l-1} \sum_{t=1}^{l-1} \lambda_t$). A positive δ_λ indicates that the current worker has higher priority than the historical average, and vice versa.

Based on the contextual information δ_K and δ_λ at each stage of worker selection, we examine all available workers at the current time period and decide whether to allocate budget 1 to selecting any worker. Intuitively, when both δ_K and δ_λ are positive (i.e., the budget is underutilized and a worker has higher priority), the selection of the considered worker is favored. When both are negative, it may not be worthwhile to spend the budget. The other cases, when one is positive and the other is negative, are more complex because we would like to spend budget on workers with higher priority but also need to save budget for future time periods in case better worker candidates arrive.

Our solution to the sequential decision problem is inspired by the well-known multiarmed bandit problem (MAB), which has been widely studied and applied to decisions in clinical trials, online news recommendation, and portfolio design. ϵ -greedy, which achieves a tradeoff between exploitation and exploration, proves to be often hard to beat by other MAB algorithms (Vermorel and Mohri 2005). Hence, we propose an adaptive budget allocation strategy based on the *contextual ϵ -greedy* algorithm (Li et al. 2010). We illustrate our solution in Figure 4.

At each stage of the local heuristic, a binary decision to make is whether to allocate budget 1 to activate the current worker with the highest priority. The contextual ϵ -greedy algorithm allows us to specify an exploration-exploitation ratio (i.e., ϵ) based on the worker's context (i.e., δ_K and δ_λ). As depicted in Figure 4, an ϵ_i -greedy algorithm is used to determine whether to select the current worker based on his δ_K and δ_λ . For each case, a YES decision is made with $1 - \epsilon_i$ probability and a NO decision with ϵ_i probability. By default, we set $\epsilon_1 = 1$ and $\epsilon_4 = 0$ to reflect NO and YES decisions, respectively, as discussed earlier. When δ_K and δ_λ have different signs, the decision is not as straightforward as the other cases, and thus we set $\epsilon_2 = \epsilon_3 = 0.5$ to allow YES and NO decisions with equal probabilities. The pseudocode of our adaptive algorithm is depicted in Algorithm 2.

5.2.2 Historical Workload. Previously, our solution is simplified by considering $\{K^{equal}[t]\}$ as the baseline budget strategy. Since human activity exhibits temporal patterns, understanding those patterns may help to guide budget allocation. Therefore, we propose to compute a baseline budget strategy with historical data that captures the expected worker/task patterns. The study in Musthag and Ganesan (2013) shows the time-of-day usage patterns of workers in mobile crowdsourcing applications. The activity peaks are between 4 and 7 pm when workers leave their day jobs. Similar patterns are observed in Foursquare and Gowalla datasets in Figure 5. Figure 5(a)

Fig. 5. Daily and weekly human activity patterns.

ALGORITHM 2: ADAPTIVE BUDGET ALGORITHM (ADAPT)

```

1: Input:  $W_i, T_i$ , total budgets  $K$ 
2: Output: selected workers  $L_i$ 
3: Init  $R = T_i$ ; used budget  $K_{used} = 0$ ; average gain  $\overline{\lambda}_{i-1} = 0$ 
4: Budget allocation  $K^{equal}[]$  with Equal strategy
5: For each time period  $s_i$ 
6:   Perform Lines 4–8 from Algorithm 1
7:   Remained budget  $K_i = K - K_{used}$ 
8:   If  $i = Q$ , then  $\delta_K = K_i$  {the last time period}
9:   Otherwise,  $\delta_K = (\sum_{t=1}^i K^{equal}[t]) - K_{used}$ 
10:  While  $|L_i| < K_i$  and  $R$  is not empty:
11:    Select  $w_i$  in  $W_i$  with highest  $\lambda_i$ 
12:    Delta gain  $\delta_\lambda = \lambda_i - \overline{\lambda}_{i-1}$ 
13:    If  $\delta_\lambda \leq 0$  and  $\delta_K \leq 0$  and  $rand(0, 1) \leq \epsilon_1$ , then break
14:    If  $\delta_\lambda \leq 0$  and  $\delta_K > 0$  and  $rand(0, 1) \leq \epsilon_2$ , then break
15:    If  $\delta_\lambda > 0$  and  $\delta_K > 0$  and  $rand(0, 1) \leq \epsilon_3$ , then break
16:    If  $\delta_\lambda > 0$  and  $\delta_K \leq 0$  and  $rand(0, 1) \leq \epsilon_4$ , then break
17:     $\delta_K = \delta_K - 1$ 
18:    Perform Line 11 from Algorithm 1
19:   $\overline{K_{used}} = K_{used} + |L_i|$  {update the budget}
20:   $\overline{\lambda}_i = (\overline{\lambda}_{i-1}(Q - 1) + \lambda_i)/Q$ 
21:  Perform Lines 12, 13 from Algorithm 1

```

shows that the hourly count of check-ins presents three peaks (i.e., during lunch and morning/afternoon commutes). In Figure 5(b), we can observe peak check-in activities during weekends.

With historical worker and task information, we can leverage the optimal budget allocation strategy in the recent past and use it as the baseline strategy in Equation 6. We propose to learn the budget allocation of previous time periods, namely *workload*, using the greedy algorithm for the offline *dMTC* problem. To guide future budget allocation decisions, the previous workload $K^{prev}[]$ will be used as the baseline in Equation 6. We will empirically evaluate our proposed solutions in the experiment section.

6 WORKER OVERLOAD

In this section, we present an enhancement to our solution in order to avoid repetitive activations of the same workers. The practical implication is that those workers who locate in popular areas can be repeatedly selected by our heuristics. Overloading workers may result in undesirable

consequences, such as tasks being rejected and the workers either feeling too annoyed or stressed out to report. Several recent studies (Alfarrarjeh et al. 2015; Zhang et al. 2016; Liu et al. 2016a) also discuss the issue of overassigning tasks to workers. These studies minimize worker overloading by balancing the workload of the workers. For example, the objective is to find an assignment that minimizes the variance of the workload among workers (i.e., maximize the so-called social fairness (Liu et al. 2016a)). Another work (Alfarrarjeh et al. 2015) also aims to assign a similar number of tasks to each worker. However, none of these studies considers task assignment and worker overloading as a multiobjective optimization problem.

Our idea is to minimize the phenomenon of overloading. This requires us to maintain the number of times each worker w_{id} has been activated up to time i , $map < w_{id}, count_i(w_{id}) >$. The counter is defined as:

$$count_i(w_{id}) = \sum_{k=1}^i \sum_{j=1}^{|W_k|} d(w_k^j) [w_k^j.id = id], \quad (8)$$

where $d(w_k^j)$ represents a decision to select the j^{th} worker at time k : $d(w_k^j) = 1$ if the worker is selected; otherwise $d(w_k^j) = 0$. The brackets enclose a condition that includes the term $d(w_k^j)$ to the sum if and only if w_k^j is identified by the same id .

We include minimization of worker overloading as another objective to coverage maximization. In the following, we formulate a multiobjective optimization (MOO) problem and propose solutions in both fixed-budget and dynamic-budget scenarios.

6.1 Fixed Budget $fMTC$

In the fixed-budget setting, we formally define the MOO problem for each time instance i :

$$\text{Maximize } \bigcup_{w_i^j \in W_i} C(w_i^j) \quad (9a)$$

$$\text{Minimize } \max_{w_i^j \in W_i} (count_i(w_i^j)) \quad (9b)$$

$$\text{s.t. } \sum_{j=1}^{|W_i|} d(w_i^j) \leq K_i. \quad (9c)$$

Equation (9a) maximizes the coverage of the selected workers while Equation (9b) minimizes the highest activation count across all workers present at the time. The constraint in Equation (9c) ensures that the number of selected workers does not exceed the budget K_i at each time instance.

Rather than coming up with heuristics to sort the workers according to two objectives, we adopt a widely used approach, the Nondominated Sorting Genetic Algorithm (NSGA) [Srinivas and Deb 1994], to solve the MOO formulation for each time instance. Intuitively, nondominated sorting is used to maintain stable nondominated fronts (i.e., subpopulations of good individuals) in a multidimensional space, where each dimension corresponds to an objective. A nondominated front, also referred to as a Pareto optimal, is a solution where none of the objective functions can be improved in value without degrading other objective values. The advantage of genetic algorithms is that they simultaneously deal with a set of possible solutions (i.e., population), which enables us to find several members of the Pareto optimal set in a single run of the algorithm. We outline our solution based on the NSGA algorithm in Algorithm 3.

ALGORITHM 3: NSGA ALGORITHM

```

1: Population  $P(t) \leftarrow \text{RandomInit}$ ,  $t \leftarrow 0$ 
2: While  $t < \text{maxgen}$ 
3:    $P(t)' \leftarrow \text{Select nondominated fronts } \{P(t)\}$ , ranked by Equations (9a) and (9b)
4:    $P(t)'' \leftarrow \text{Mutation } \{P(t)' \cup \text{Crossover } \{P(t)'\}\}$ 
5:    $P(t+1) \leftarrow P(t) \cup P(t)''$ 
6:    $t \leftarrow t + 1$ 
7: Select best solution from  $P(t)$ , ranked by Equation (10)

```

The results of NSGA, at the end of while loop, include a set of nondominated fronts. Subsequently, in Line 7, we select the best individual solution based on a weighted sum of objective values:

$$\alpha \left| \bigcup C(w_i^j) \right| / |T_i| - (1 - \alpha) \max(\text{count}(w_i^j)) / Q. \quad (10)$$

In Equation (10), α is a linear coefficient, $0 < \alpha < 1$, to specify the weight for each objective. The higher α , the more important the objective of Equation (9a) in comparison to that of Equation (9b). The minus sign indicates the minimization objective in Equation (9b). Both objective functions are normalized by the total number of tasks $|T_i|$ and the total number of time instances Q , respectively. In our experiments, we adopted the NSGA-II version (Deb et al. 2002) for implementing Algorithm 3 and set $\text{maxgen} = 50,000$.

6.2 Dynamic Budget dMTC

In the dynamic-budget setting, the MOO formulation is similar to Equation (9a, 9b), but the total budget is constrained over all time periods. Therefore, the constraint in Equation (9c) is replaced by the following constraint:

$$\sum_i \sum_{j=1}^{|W_i|} d(w_i^j) \leq K. \quad (11)$$

In the online setting, we need to simultaneously consider the task coverage and the number of activations of the candidate worker in order to optimize both objectives. As a result, we modify the adaptive strategy in Section 5.2.1 and define the gain λ_l of the current worker w_i^j in Equation (7) to be a linear combination of the number of previous activations and his priority:

$$\lambda_l = \alpha \cdot \text{priority}(w_i^j) / |T_i| - (1 - \alpha) \text{count}(w_i^j) / Q. \quad (12)$$

In Equation (12), $\text{priority}(w_i^j)$ and $\text{count}(w_i^j)$ are respectively the priority of the worker calculated by any previously proposed local heuristic and the number of times that worker was selected. The coefficient α can be varied to balance the importance of overloading and priority.

7 DISTANCE-BASED TASK UTILITY

Thus far, our goal is to maximize the number of assigned tasks, assuming that assigning to any worker within the task region is equivalent. However, in practice, the assignment of a task to a nearby worker may yield higher *utility* than that of a farther worker (Miao et al. 2016). Thus, in this section, we aim to generalize the binary-utility model to a distance-based-utility variant (i.e., maximizing the *utility* of covered tasks). We assume the utility of worker w 's response to task t is a function of the spatial distance between them: $\text{utility}(t, w) = f(\text{dist}(t, w))$. And f is a decreasing function of worker-task distance. Intuitively, the utility is at its highest when the worker is co-located with the task, and it decreases as the worker-task distance increases. The utility is zero if the distance is larger than the task radius. We consider three cases depicted in

Fig. 6. Distance-based utility functions.

Table 2. Summaries of Real-World Datasets

Name	#Tasks	#Workers	MTD ¹¹	$ s_i $
Foursquare	89,968	45,138 (90/km ²)	16.6km	1 hour
Gowalla	151,075	6,160 (35/km ²)	3.6km	1 day

Figure 6: (i) *Binary*, where utility has value 1/0; (ii) *Linear*, where utility decreases linearly with the worker–task distance; and (iii) *Zipf*, where utility follows *Zipfian* distribution with skewness parameter s . The higher the value of s , the faster utility drops.

This extension can be incorporated into all the previously developed algorithms. Specifically, Algorithm 1 (Line 10) now chooses the worker who maximizes utility increase at each stage:

$$priority(w_i^j) = \sum_{t_i^k \in C(w_i^j) \cap R} f(dist(t_i^k, w_i^j)). \quad (13)$$

With the temporal heuristic, Equation (3) becomes:

$$priority(w_i^j) = \sum_{t_i^k \in C(w_i^j) \cap R} \frac{f(dist(t_i^k, w_i^j))}{t_i^k \cdot (\delta + s) - i}. \quad (14)$$

In the same fashion, with the spatial heuristic, Equation (5) becomes:

$$priority(w_i^j) = \sum_{t_i^k \in C(w_i^j) \cap R} \frac{f(dist(t_i^k, w_i^j))}{1 + RE(t_i^k)}. \quad (15)$$

With the adaptive budget strategies, the gain of a candidate worker is adapted similarly.

8 PERFORMANCE EVALUATION

8.1 Experimental Methodology

We adopted real-world datasets from location-based applications, summarized in Table 2, to emulate SC workers and tasks. We consider Gowalla and Foursquare users as SC workers and the venues as tasks. The Gowalla dataset contains check-ins for 224 days in 2010, including more than 100,000 spots (e.g., restaurants) within the state of California. By considering each day as a unit time period, all the users who checked in during a day are available workers for that time period in our setting. The Foursquare dataset contains the check-in history of 45,138 users to 89,968 venues over 384 hours in Pittsburgh, Pennsylvania. We considered each hour as a unit time period for this dataset.

¹¹MTD: Mean Travel Distance (To et al. 2014).

We generated a range of datasets with the SCAWG toolbox (To et al. 2016a) by utilizing real-world worker-task spatial distributions and varying their arrival rate. We generated a worker count following COSINE (default) and POISSON distributions with mean = 50 and set the default value of task count per time period to be constant (i.e., 1000). We denote Go-POISSON a dataset that uses Gowalla for the spatial distribution and POISSON for the worker arrival rate.

In all of our experiments, for the Gowalla dataset, we varied the total number of time periods $Q \in \{7, 14, 28, 56\}$ and the budget $K \in \{56, 112, 224, 448, 896, 1288\}$. For Foursquare, $Q \in \{24, 48, 72, 96\}$ and $K \in \{24, 48, 96, 192, \dots, 1536\}$ because we modeled a time period as one hour. The task duration δ was randomly chosen from 1 to Q , and the task radius $r \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ km. The choices of r and δ values are defined by the CHRS experts. Default values are shown in boldface. Finally, experiments were run on an Intel(R) Core(TM)i7-2600 CPU at 3.40GHz with 8GB of RAM.

8.2 Experimental Results

In the following, we evaluate our solutions in terms of the number of covered tasks (i.e., *task coverage*). We first show the performance of offline solutions (Section 8.2.1). Then we present the results for the online scenario, including local heuristics, adaptive budget strategy, and workload heuristic (Section 8.2.2). We next show the results of distance-based utility and worker overloading (Section 8.2.3), followed by runtime measurements (Section 8.2.4).

8.2.1 Optimal Solutions for Offline Setting. We implemented the offline solutions for the two problem variants, *fMTC* and *dMTC* (Section 4), using integer linear programming. These algorithms provide optimal results that are used as the upper bounds of the online algorithms. Figure 7(a) illustrates the results for Go-POISSON by varying the total budget K . As expected, higher budget yields higher coverage. Also, *DynamicOff* yields higher coverage than *FixedOff* as *FixedOff* is the special case of *DynamicOff*. However, the higher the budget, the smaller the performance gap between them. This effect can be explained by the diminishing return property of the max cover problem. That is, the more workers are selected, the smaller gain of each selected worker. We also evaluated the offline solutions by varying task radius r (Figure 7(b)). Intuitively, when r increases, more workers are located within the task's spatial range. This means that a worker is eligible to perform more tasks, which yields higher task coverage.

Similar trends were observed for Go-POISSON, as shown in Figures 7(c) and 7(d). We observe a small difference between *FixedOff* and *DynamicOff* for Go-POISSON in Figure 7(d). However, when the arrival rate has high variance, such as in Go-COSINE, *DynamicOff* shows more coverage over *FixedOff* in Figure 7(b). The reason is that *FixedOff* uses a constant budget to the time periods with high spikes while *DynamicOff* can allocate more budget to those time periods to cover more tasks.

8.2.2 Local Heuristics and Budget Allocations for Online Setting.

The Performance of Heuristics: We evaluate the performance of the online heuristics for *fMTC* from Section 5.1, *Basic*, *Spatial* and *Temporal*. Figures 8(a) shows the improvements of *Spatial* and *Temporal* over *Basic* on Go-COSINE. When the budget is high, we observe that the simple heuristic *Basic* already obtains results close to the optimal solution. This is because most workers are selected. Furthermore, *Spatial* and *Temporal* yield 2% and 4% higher coverage than *Basic* (at $K = 224$), and their performance converges as K increases. Similar trends can be observed when increasing the task radius r in Figure 8(b).

Adaptive Budget Allocation Strategy: We evaluate the performance of the adaptive budget allocation strategy in Section 5.2.1 by comparing it with three baseline strategies inspired by a few previous studies: namely, *Equal* (To et al. 2016b; Tran-Thanh et al. 2013), *Random* (Tran-Thanh et al. 2013), and *Naive* (Kazemi et al. 2013; Ji et al. 2016; Zhang et al. 2015). In the *Equal* budget strategy,

Fig. 7. Performance of offline solutions with Go-COSINE and Go-POISSON.

Fig. 8. Performance of local heuristics in the fixed-budget scenario.

the total budget is allocated equally to the time intervals. In the *Random* budget strategy, a random positive number k_i is generated for each time interval i , and then each time interval i is given a budget $K_i = K * \frac{k_i}{\sum_1^Q k_i}$. In the *Naive* budget strategy, there is no particular limitation for the budget of each time interval; the budget is used until no more workers are available or entire budget is exhausted. We use local heuristic *Basic* to compare the performances of budget strategies. Figure 9 shows the results of the budget strategies when varying the total budget K . *AdaptB* is shown to be the best in coverage. *EqualB* and *RandomB* do not perform as well as *AdaptB* as they lack an intelligent budget allocation strategy. *NaiveB* performs poorly as it selects workers on a first-come-first-served basis without considering future time intervals. As a result, the total budget is quickly exhausted during the first few time intervals. The difference between *AdaptB* and the others is

Fig. 9. Performance of adaptive budget allocation strategies.

Fig. 10. Performance of adaptive budget allocation in the dynamic-budget scenario.

higher with GO-COSINE because it has more fluctuation in the number of workers over time. This shows that *AdaptB* can quickly adapt to the dynamic arrivals of workers.

Historical Workload Improvement: We also evaluate the performance of the adaptive budget allocation strategy applied with local heuristics, in which *Temporal* is shown to perform better than *Basic*. Figure 10 shows the results of *EqualB*, *AdaptB*, *AdaptT*, and *AdaptTW* (*AdaptT* with historical workload improvement) when varying total budget K and task radius r on Go-COSINE and Go-POISSON datasets. We include *DynamicOff* as the optimal result for reference. As can be seen in that figure, with small budgets, *AdaptTW*, which uses historical optimal workload as the baseline budget strategy, has higher coverage than the others, and, with $K = 56$, *EqualB* performs better than *AdaptB*, *AdaptT*. The reason is that, with small budgets, *AdaptB* and *AdaptT* do not

Fig. 11. Boxplots for various budget allocation strategies.

Fig. 12. Performance of *AdaptT* on Go-COSINE with distance-based utility. The overlapping ratio indicates the percentages of workers who are selected in both binary and corresponding distance-based utility model.

have enough contextual information. With a higher budget ($K \geq 448$), the adaptive algorithms perform better than *EqualB*, and their results are close to the optimal result. We observe similar results when varying r .

We further study the performance of various budget allocation strategies by plotting task coverage across multiple workloads using boxplots. Figure 11 shows the results of the techniques with the default parameter setting. As can be seen in the figure, the adaptive algorithms perform better than *EqualB*, especially with the GO-COSINE dataset. While *AdaptTW* has the highest median, minimum, and maximum values, *AdaptT* is the most stable method with the smallest difference between the minimum and maximum values.

8.2.3 Distance-Based Task Utility and Worker Overload.

Distance-Based Task Utility: We show the performance of *AdaptT* under distance-based functions for task utility from Section 7. In Figure 12(a), we observe that the obtained task coverage in the cases of Linear and Zipfian are similar to the result in the binary-utility model. We also present the overlapping ratio of selected workers between distance-based utility and binary utility in Figure 12(b). The figure shows that when total budget increases, initially, the ratio decreases because more and different workers can be selected; then ratio increases because, when the budget is large enough, the most workers are selected in both cases, resulting in large overlaps.

Worker Overload Minimization: In this section, we evaluate the performance of MOO techniques in both fixed-budget and dynamic-budget scenarios. The techniques are evaluated in terms of balancing the tradeoff between maximizing task coverage and minimizing worker overload. In

Fig. 13. Performance of *EqualGA* and *AdaptT-MOO* when varying α .

the fixed-budget scenario (Section 6.1), *EqualGA* refers to the equal-budget strategy with NSGA in Algorithm 3. We observe that varying coefficient α does not significantly change task coverage (Figure 13(a)). This is due to the equal allocation of the total budget to each time period, which yields suboptimal task coverage. By increasing α , the average number of activations per worker in Figure 13(b) shows a slightly decreasing trend due to a higher weight on the second objective. In the dynamic-budget scenario, *AdaptT-MOO* refers to adaptive budget allocation with a temporal local heuristic and MOO in Section 6.2. Figure 13(c) shows that the task coverage is quite stable as α increases while the average number of activation decreases significantly. This means that our adaptive budget allocation strategy achieves workload balancing among the workers at a very small cost in utility. Furthermore, without loss of generality, based on the observations, we set $\alpha = 0.1$ for the following experiments.

Figure 14 shows the distribution of activation counts of selected workers when $K = 448$. In the fixed-budget setting, it is shown that *EqualGA* does not cover as many tasks as *EqualB*, but it activates more workers for a small number of times (i.e., 1, 2, and 3 times). In the dynamic-budget setting, *AdaptT-MOO* also has more workers with a small number of activations and yields comparable task coverage compared to *AdaptT*. We conclude that our solutions can mitigate worker overloading without compromising task assignment.

8.2.4 Runtime Measurements. Figure 15 shows the runtime performance of our online algorithms by varying the number of tasks per time period. As expected, the runtime linearly increases when the number of tasks grows. In the fixed-budget scenario, the runtimes of the local heuristics (e.g., *Temporal*) are the same as *Basic* while the runtime of *EqualGA* is higher due to having a large number of iterations for Algorithm 3. We do not show the runtime of *Spatial* heuristic and *Zipfian* utility model, but their runtimes are similar to *EqualB* and *EqualT-Linear*, respectively.

Fig. 14. Worker activation count distribution of MOO-based algorithms in the fixed and dynamic budget scenarios ($K = 448$, $\alpha = 0.1$).

Fig. 15. Average runtime per time period with Go-COSINE and Fo-COSINE.

In the dynamic-budget scenario, the runtime of *AdaptTW* is higher than *AdaptT*, *AdaptT-Linear*, and *AdaptT-MOO*. This suggests that the workload heuristic significantly increases the overhead of *AdaptT*. However, the MOO extension does not incur observable runtime overhead in the dynamic-budget scenario.

9 DISCUSSION

Existing studies show that knowing the worker mobility pattern *a priori* can improve the efficiency of task assignment (Ji et al. 2016; Zhang et al. 2015). Our solution does not consider individual worker mobility patterns (i.e., the worker's trajectory) for task assignment. However, our heuristics (Section 5.1.3) do consider worker population mobility patterns by prioritizing tasks whose

locations are not likely to be visited by many workers in the future. Furthermore, our dynamic budget algorithm (Algorithm 2) takes into account the dynamic arrivals of workers and tasks as well as their co-location relationship.

It is worth noting that in our problem settings (i) task assignment is real-time and online, and (ii) workers are not required to travel to perform tasks. Workers can respond to a task immediately after receiving the task notifications from the SC-server. Therefore, they do not need to perform a sequence of tasks as in a typical mobile crowdsourcing application, where workers often chain multiple tasks to maximize their earnings while minimizing travel time (Ji et al. 2016). In addition, workers' trajectories within the same task region would not have much impact in our problem setting as the workers are not required to travel to perform the task. Obviously, as the workers move, they may become relevant to another spatial task and/or irrelevant to the prior task, which can be represented as the addition and deletion of a worker in our framework at a given snapshot.

10 CONCLUSION

Motivated by weather crowdsourcing applications, we introduced the problem of Hyperlocal SC, where tasks can be performed by workers within their spatiotemporal vicinity. We studied task assignment in Hyperlocal SC to maximize the covered tasks without exceeding the budget for activating workers. A range of problem variants was considered, including offline versus online, budget constraint for each time period versus for the entire campaign, single objective versus multiple objectives, and binary versus distance-based utility. We showed that the offline variants are NP-hard, and we proposed several local heuristics and a dynamic budget allocation for the online scenario, which utilize the spatial and temporal properties of workers and tasks. We generated SC workloads with the SCAWG tool and conducted extensive experiments. We concluded that *AdaptT*, which merits the temporal local heuristic and dynamic budget allocation, is the superior technique in terms of utility and runtime. The extensions to measure distance-based utility and to minimize worker overloading were shown to be very effective and do not impose significant runtime overhead. As future work, we will consider the nonuniform activation cost of the workers, which represents the reputation or the compensation demand of each worker. We will also consider assigning a task to multiple workers to improve the quality of collected data and utilizing known worker mobility patterns to boost task assignment.

ACKNOWLEDGMENTS

We would like to thank CHRS researchers, especially Dr. Phu Dinh Nguyen, for leading the development of the iRain project: <http://irain.eng.uci.edu/>.

REFERENCES

- 2016. iRain: new mobile App to promote citizen-science and support water management. Retrieved from <http://en.unesco.org/news/irain-new-mobile-app-promote-citizen-science-and-support-water-management>.
- Abdullah Alfarrarjeh, Tobias Emrich, and Cyrus Shahabi. 2015. Scalable spatial crowdsourcing: A study of distributed algorithms. In *Proceedings of the 2015 16th IEEE International Conference on Mobile Data Management (MDM)*, Vol. 1. IEEE, 134–144.
- Chandra Chekuri and Amit Kumar. 2004. Maximum coverage problem with group budget constraints and applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Springer, 72–83.
- Peng Cheng, Xiang Lian, Lei Chen, Jinsong Han, and Jizhong Zhao. 2016. Task assignment on multi-skill oriented spatial crowdsourcing. *IEEE Transactions on Knowledge and Data Engineering* 28, 8 (2016), 2201–2215. arxiv:1510.03149
- Peng Cheng, Xiang Lian, Zhao Chen, Rui Fu, Lei Chen, Jinsong Han, and Jizhong Zhao. 2015. Reliable diversity-based spatial crowdsourcing by moving workers. *Proceedings of the VLDB Endowment* 8, 10 (2015), 1022–1033. arxiv:arXiv:1412.0223v1
- Justin Cranshaw, Eran Toch, Jason Hong, Aniket Kittur, and Norman Sadeh. 2010. Bridging the gap between physical location and online social networks. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing*. ACM.

- Hung Dang, Tuan Nguyen, and Hien To. 2013. Maximum complex task assignment: Towards tasks correlation in spatial crowdsourcing. In *Proceedings of International Conference on Information Integration and Web-based Applications & Services*. ACM, 77.
- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation* 6, 2 (2002), 182–197.
- Dingxiong Deng, Cyrus Shahabi, Ugur Demiryurek, and Linhong Zhu. 2016. Task selection in spatial crowdsourcing from worker's perspective. *Geoinformatica* 20, 3 (JJul 2016), 529–568.
- Bruce Dorminey. 2014. Crowdsourcing the weather. Retrieved from <http://www.forbes.com/sites/brucedorminey/2014/02/26/crowdsourcing-as-the-future-of-weather-forecasting/>.
- Uriel Feige. 1998. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)* 45, 4 (1998), 634–652.
- Dawei Gao, Yongxin Tong, Jieying She, Tianshu Song, Lei Chen, and Ke Xu. 2017. Top-k team recommendation and its variants in spatial crowdsourcing. *Data Science and Engineering* (2017), 1–15.
- Hui Gao, Chi Harold Liu, Wendong Wang, Jianxin Zhao, Zheng Song, Xin Su, Jon Crowcroft, and Kin K Leung. 2015. A survey of incentive mechanisms for participatory sensing. *IEEE Communications Surveys & Tutorials* 17, 2 (2015), 918–943.
- Bin Guo, Yan Liu, Wenle Wu, Zhiwen Yu, and Qi Han. 2016. ActiveCrowd: A framework for optimized multitask allocation in mobile crowdsensing systems. *IEEE Transactions on Human-Machine Systems* 47, 3 (2016), 392–403.
- Shibo He, Dong-Hoon Shin, Junshan Zhang, and Jiming Chen. 2014. Toward optimal allocation of location dependent tasks in crowdsensing. In *Proceedings of the 2014 IEEE INFOCOM*. IEEE, 745–753.
- Dorit S. Hochbaum. 1996. Approximating covering and packing problems: Set cover, vertex cover, independent set, and related problems. In *Approximation Algorithms for NP-hard Problems*. PWS Publishing Co., 94–143.
- Huiqi Hu, Yudian Zheng, Zhifeng Bao, Guoliang Li, Jianhua Feng, and Reynold Cheng. 2016. Crowdsourced POI labelling: Location-aware result inference and task assignment. In *Proceedings of the 2016 IEEE 32nd International Conference on Data Engineering (ICDE 2016)*. IEEE, 61–72.
- Sheng Gong Ji, Yu Zheng, and Tianrui Li. 2016. Urban sensing based on human mobility. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 1040–1051.
- Haiming Jin, Lu Su, Bolin Ding, Klara Nahrstedt, and Nikita Borisov. 2016. Enabling privacy-preserving incentives for mobile crowd sensing systems. In *Proceedings of the 2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 344–353.
- Thivya Kandappu, Archan Misra, Shih-fen Cheng, Nikita Jaiman, and Randy Tandriansyah. 2016. Campus-scale mobile crowd-tasking : Deployment & behavioral insights. In *Proceedings of the 19th ACM Conference on Computing and Cooperative Work and Social Computing*. ACM Press, New York, 798–810.
- Leyla Kazemi and Cyrus Shahabi. 2012. GeoCrowd: Enabling query answering with spatial crowdsourcing. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems*.
- Leyla Kazemi, Cyrus Shahabi, and Lei Chen. 2013. GeoTruCrowd: Trustworthy query answering with spatial crowdsourcing. In *Proceedings of the 21st ACM SIGSPATIAL GIS 2013*.
- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 661–670.
- Yu Li, Man Lung Yiu, and Wenjian Xu. 2015. Oriented online route recommendation for spatial crowdsourcing task workers. In *Lecture Notes in Computer Science (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, Vol. 9239. Springer International Publishing, 137–156.
- Qing Liu, Talel Abdesslem, Huayu Wu, Zihong Yuan, and Stéphane Bressan. 2016a. Cost minimization and social fairness for spatial crowdsourcing tasks. In *International Conference on Database Systems for Advanced Applications*. Springer, 3–17.
- Yan Liu, Bin Guo, Yang Wang, Wenle Wu, Zhiwen Yu, and Daqing Zhang. 2016b. TaskMe: Multi-task allocation in mobile crowd sensing. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive Ubiquitous Computing (UbiComp'16)*. ACM Press, New York, New York, 403–414.
- Chunyan Miao, Han Yu, Zhiqi Shen, and Cyril Leung. 2016. Balancing quality and budget considerations in mobile crowdsourcing. *Decision Support Systems* 90 (2016), 56–64.
- Mohamed Musthag and Deepak Ganesan. 2013. Labor dynamics in a mobile micro-task market. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 641–650.
- Layla Pournajaf, Li Xiong, Vaidy Sunderam, and Slawomir Goryczka. 2014. Spatial task assignment for crowd sensing with cloaked locations. In *Proceedings of the IEEE International Conference on Mobile Data Management*, Vol. 1. IEEE, 73–82.
- André Sales Fonteles, Sylvain Bouveret, and Jérôme Gensel. 2016. Trajectory recommendation for task accomplishment in crowdsourcing a model to favour different actors. *Journal of Location Based Services* 10, 2 (apr 2016), 125–141.
- Zheng Song, Chi Harold Liu, Jie Wu, Jian Ma, and Wendong Wang. 2014. QoI-aware multitask-oriented dynamic participant selection with budget constraints. *IEEE Transactions on Vehicular Technology* 63, 9 (2014), 4618–4632.

- Nidamarthi Srinivas and Kalyanmoy Deb. 1994. Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation* 2, 3 (1994), 221–248.
- Hien To, Mohammad Asghari, Dingxiong Deng, and Cyrus Shahabi. 2016a. SCAWG: A toolbox for generating synthetic workload for spatial crowdsourcing. In *Proceedings of the 2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 1–6.
- Hien To, Liyue Fan, Luan Tran, and Cyrus Shahabi. 2016b. Real-time task assignment in hyperlocal spatial crowdsourcing under budget constraints. In *Proceedings of the 2016 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 1–8.
- Hien To, Gabriel Ghinita, Liyue Fan, and Cyrus Shahabi. 2017. Differentially private location protection for worker datasets in spatial crowdsourcing. *IEEE Transactions on Mobile Computing* 16, 4 (2017), 934–949.
- Hien To, Gabriel Ghinita, and Cyrus Shahabi. 2014. A framework for protecting worker location privacy in spatial crowdsourcing. *Proceedings of the VLDB Endowment* 7, 10 (2014), 919–930.
- Hien To, Cyrus Shahabi, and Leyla Kazemi. 2015. A server-assigned spatial crowdsourcing framework. *ACM Transactions on Spatial Algorithms and Systems* 1, 1 (2015), 2.
- Yongxin Tong, Jieying She, Bolin Ding, Libin Wang, and Lei Chen. 2016. Online mobile micro-task allocation in spatial crowdsourcing. In *Proceedings of the 2016 IEEE 32nd International Conference on Data Engineering (ICDE 2016)*. IEEE, 49–60.
- Long Tran-Thanh, Matteo Venanzi, Alex Rogers, and Nicholas R. Jennings. 2013. Efficient budget allocation with accuracy guarantees for crowdsourcing classification tasks. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*. 901–908.
- Umar ul Hassan and Edward Curry. 2014. A multi-armed bandit approach to online spatial task assignment. In *Proceedings of the 11th IEEE International Conference on Ubiquitous Intelligence and Computing (UIC)*.
- Matteo Venanzi, Alex Rogers, and Nicholas R. Jennings. 2013. Crowdsourcing spatial phenomena using trust-based heteroskedastic gaussian processes. In *Proceedings of the 1st AAAI Conference on Human Computation and Crowdsourcing*.
- Joannes Vermorel and Mehryar Mohri. 2005. Multi-armed bandit algorithms and empirical evaluation. In *Machine Learning: ECML 2005*. Springer, 437–448.
- Leye Wang, Dingqi Yang, Xiao Han, Tianben Wang, Daqing Zhang, and Xiaojuan Ma. 2017. Location privacy-preserving task allocation for mobile crowdsensing with differential geo-obfuscation. In *Proceedings of the 26th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 627–636.
- Mingjun Xiao, Jie Wu, Liusheng Huang, Yunsheng Wang, and Cong Liu. 2015. Multi-task assignment for crowdsensing in mobile social networks. In *Proceedings of the 2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE, 2227–2235.
- Bo Zhang, Zheng Song, Chi Harold Liu, Jian Ma, and Wendong Wang. 2015. An event-driven qoi-aware participatory sensing framework with energy and budget constraints. *ACM Transactions on Intelligent Systems and Technology (TIST)* 6, 3 (2015), 42.
- Daqing Zhang, Haoyi Xiong, Leye Wang, and Guanling Chen. 2014. Crowdrecruiter: Selecting participants for piggyback crowdsensing under probabilistic coverage constraint. In *ACM UbiCom 2016*. ACM, 703–714.
- Hongli Zhang, Zhikai Xu, Xiaojiang Du, Zhigang Zhou, and Jiantao Shi. 2016. CAPR: Context-aware participant recruitment mechanism in mobile crowdsourcing. *Wireless Communications and Mobile Computing* 16, 15 (2016), 2179–2193.
- Xinglin Zhang, Zheng Yang, Yunhao Liu, Jianqiang Li, and Zhong Ming. 2017. Toward efficient mechanisms for mobile crowdsensing. *IEEE Transactions on Vehicular Technology* 66, 2 (2017), 1760–1771.
- Yu Zheng, Licia Capra, Ouri Wolfson, and Hai Yang. 2014. Urban computing: Concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 3 (2014), 38.

Received December 2016; revised March 2017; accepted March 2017