

Incentivizing Quality-based Data Crowdsourcing

Xiaowen Gong
Auburn University
xgong@auburn.edu

Abstract

Data quality is a key metric of data crowdsourcing, which captures how accurate the data is compared to the ground truth. In general, workers participating in a crowdsourcing task have diverse intrinsic quality, depending on their characteristics (e.g., location, expertise). To improve data quality, it is important to incentivize workers to truthfully make effort in the task and report their data. In this paper, we study effort elicitation in data crowdsourcing where workers have diverse quality and know their quality, while it is unknown to the requester. We show that for one-stage effort elicitation, a symmetric quality-based threshold strategy is an equilibrium for the workers. Then we propose an online quality learning mechanism with sequential effort elicitation. We show that under this mechanism, a symmetric quality-based threshold strategy is an approximate equilibrium, where the approximate error goes to 0 asymptotically. Based on the preliminary results in this paper, several extensions and generalizations of the problem here will be studied in future work.

1 Introduction

Data crowdsourcing (referred to as “crowdsourcing” for brevity) has found a wide range of applications. Typical applications involves physical sensing tasks (also known as “crowdsensing”) such as spectrum sensing, traffic monitoring, and environmental monitoring. In principle, crowdsourcing leverages the “wisdom” of a potentially large crowd of workers (e.g., mobile users) for a crowdsourcing task. A key advantage of crowdsourcing lies in that it can exploit the diversity of inherently inaccurate data from many workers by aggregating the data obtained by the crowd, such that the data accuracy (also referred to as “data quality”) after aggregation can be substantially enhanced. With enormous opportunities brought by big data, crowdsourcing serves as an important first step for data mining tools to harness the power of big data in many application domains.

To fully exploit the potential of crowdsourcing, it is important to assign tasks to workers based on their quality.

A worker’s quality¹ captures the *intrinsic* accuracy of the worker’s data relative to the ground truth of the interested variable, and it generally varies for different workers depending on a worker’s characteristics (e.g., location, capabilities of sensors, expertise in a domain). For example, if the task is to detect whether a licensed frequency band is idle or occupied by a licensed user (for opportunistic spectrum access by unlicensed users), then the quality of a worker’s data is the probability of correct detection, which depends on the worker’s location relative to the licensed user. Workers generally have *diverse* quality. A worker can learn its quality based on the knowledge of its characteristics, such as its location. However, the quality of a worker’s can be its private information, which is unknown to and cannot be verified by the crowdsourcing requester. For example, a worker’s location is often its private information that is unknown to the requester.

In addition to the worker’s intrinsic quality, the data quality is also affected by the worker’s effort exerted in the task. The data quality when the worker makes effort is higher than when it makes no effort. For example, to detect whether a licensed frequency band is idle, a worker should measure the signal in that band to make an estimate, rather than making a guess without any measuring. However, a worker’s effort can also be its hidden action that cannot be observed by the requester. Due to the inaccurate nature of the data, a strategic worker may report some arbitrary data to the requester without making effort in the task, while the requester is not able to verify whether effort was actually made. Furthermore, the data itself obtained by a worker from the task could also be its private information that it can manipulate in favor of itself.

In this paper, we study effort elicitation² when they have diverse quality and know their quality, while the requester does not know workers’ quality and cannot explicitly elicit this information from the workers (as done by the truthful quality elicitation mechanisms proposed in [Gong and Shroff, 2017; 2018]). The goal of the requester is to incentivize workers to truthfully make effort in the task and report their data. In some situations, the requester performs task allocation and data aggregation with effort elicitation while being oblivious

¹We use “worker quality” and “quality” exchangeably in this paper. “Worker quality” should be distinguished from “data quality”.

²In this paper, we use “effort elicitation” instead of “effort and data elicitation” for brevity.

of workers’ quality. In other situations, the requester can learn workers’ quality on the fly while performing task allocation and data aggregation based on the learned quality, which can improve system efficiency. In this setting, the requester’s online quality learning and effort elicitation present non-trivial coupling with each other: effort elicitation affects workers’ quality and thus the quality that is learned, while online quality learning affects task allocation and thus workers’ reward which incentivizes their effort. This coupling is further complicated by the strategic interaction among workers. The analysis of online quality learning with effort elicitation draws on synergistic integration of models and methods from reinforcement learning and mechanism design.

The main contributions of this paper is to take initial steps towards the direction described above and present some preliminary results. We show that for one-stage effort elicitation, a symmetric quality-based threshold strategy is a Bayesian Nash equilibrium (BNE) for the workers. Then we propose an online quality learning mechanism with sequential effort elicitation. We show that under this mechanism, a symmetric quality-based threshold strategy is an approximate BNE, where the approximate error is upper bounded by a constant as the time horizon goes to infinity, and thus the time-averaged approximate error goes to 0 asymptotically.

2 Related Work

Incentive mechanisms for data crowdsourcing. There have been a lot of recent research on incentive mechanisms for data crowdsourcing [Duan *et al.*, 2012; Yang *et al.*, 2012; Koutsopoulos, 2013; Feng *et al.*, 2014; Tarable *et al.*, 2015; Shah and Zhou, 2015; Luo *et al.*, 2015; Wang *et al.*, 2016; Pu *et al.*, 2016; Zhang *et al.*, 2016]. Many of these mechanisms incentivize workers to truthfully reveal their participating cost of crowdsourcing, which is a strategic worker’s private information. Some recent works have studied effort and/or data elicitation in crowdsourcing [Dasgupta and Ghosh, 2013; Luo *et al.*, 2015; Jin *et al.*, 2017b; Liu and Chen, 2016b; 2016a; 2017], where a strategic worker may have incentive to manipulate to its own advantage. [Liu and Chen, 2016b; 2017] have designed effort elicitation mechanisms for workers with diverse cost, and studied the optimal reward for maximizing the requester’s payoff. Different from these works, we study effort elicitation for workers with diverse quality. This adds a new dimension to effort elicitation which lead to different results. [Liu and Chen, 2016a] has designed a bandit-based mechanism for sequential effort elicitation. This paper is different from [Liu and Chen, 2016a] in that workers’ quality is unknown to the requester, while it is assumed to be known in [Liu and Chen, 2016a].

Two recent studies [Gong and Shroff, 2017; 2018] have proposed a quality-aware crowdsourcing framework and devised truthful mechanisms for quality and effort elicitation. This paper is different from [Gong and Shroff, 2017; 2018], since the requester cannot explicitly elicit workers’ quality from them (e.g., so as to protect workers’ privacy). Although this paper uses online quality learning, it only provides estimates of workers’ quality, which is different from accurate quality elicited from workers as in [Gong and Shroff, 2017;

2018].

Quality-based data crowdsourcing. The quality of data is important for allocating crowdsourcing tasks to workers, and has been studied in a few works [Karger *et al.*, 2011; Koutsopoulos, 2013; Lee *et al.*, 2015; Shah and Zhou, 2015; Liu and Liu, 2015; Jin *et al.*, 2015; 2016; 2017a]. One interesting line of work [Jin *et al.*, 2015; 2016; 2017a] in this direction has studied truthful mechanisms for information quality based task allocation where workers have private participating cost. Some other works [Karger *et al.*, 2011; Lee *et al.*, 2015; Liu and Liu, 2015] have focused on learning the data quality of workers from their data. [Karger *et al.*, 2011; Lee *et al.*, 2015] have studied exploiting the correlation of workers’ data for different tasks. [Liu and Liu, 2015] has developed an online quality learning algorithm that asymptotically finds the workers with highest quality without knowing the ground truth of tasks. Different from [Liu and Liu, 2015], we study online quality learning with sequential effort elicitation, which presents intricate coupling between the two components.

3 Problem Formulation

We consider a crowdsourcing requester (also referred to as worker 0) recruiting a set of workers $\mathcal{N} \triangleq \{1, \dots, N\}$ to work on a task. For convenience, let $\mathcal{N}^+ \triangleq \mathcal{N} \cup \{0\}$. The structure and procedure of the crowdsourcing system is described in detail as follows.

Data observation. The crowdsourcing task is to observe and estimate an unknown and random variable of interest X . The interested variable X takes discrete values (e.g., the answer of a multi-choice question). For ease of exposition, we assume that X takes one of two possible values³ 0 and 1. After working on the task, each worker $i \in \mathcal{N}^+$ (i.e., including the requester) obtains random data D_i . The accuracy of the data D_i is quantified by the *correct probability* q_i , which is the probability that D_i is equal to the interested variable X , given by

$$q_i \triangleq \Pr(D_i = X) = p_i e_i + \underline{p}(1 - e_i).$$

Here the correct probability p_i depends on the *worker quality* p_i of worker i and the *effort* e_i exerted by worker i in the task, which is explained as follows.

Worker quality. Given that worker i makes effort in the task, the quality $p_i \in [0, 1]$ determines the correct probability q_i which quantifies how accurate D_i is. The quality p_i is an *intrinsic* coefficient that captures worker i ’s capability for the task. Note that a larger p_i means higher quality. The quality generally varies for different workers. We assume that each worker $i \in \mathcal{N}$ knows its own quality p_i (e.g., by learning the correct probability based on its location), but it is unknown to the requester. For ease of exposition, we assume that each worker’s quality p_i is within the range of $[\underline{p}, \bar{p}]$ which is known to the requester. We assume that each worker’s quality (including the requester’s) follows a distribution with CDF $F(p)$ and PDF $f(p)$ with mean p_0 , and these information is known to all the workers and the requester.

³The results of this paper can be fairly easily extended to the case of multiple possible values of the interested variable X .

Worker effort. The effort $e_i \in \{0, 1\}$ represents whether worker i makes effort in the task, where $e_i = 1$ and $e_i = 0$ indicate making and not making effort, respectively. If worker i makes effort, then the correct probability q_i of worker i is equal to the worker quality p_i ; otherwise, q_i is equal to \underline{p} , e.g., which means that worker i simply makes a guess of X randomly according to the prior distribution. To ensure that making effort is meaningful, we assume that $p_i > \underline{p}$. Therefore, given the quality p_i , making effort $e_i = 1$ means a larger correct probability q_i and thus higher accuracy of D_i than not making effort. The binary effort model (i.e., either making effort or not) is reasonable (also used in, e.g., [Dasgupta and Ghosh, 2013; Liu and Chen, 2016b; 2017]), as workers' behavior tend to be simple in practice. We assume that each worker i can control its effort e_i , but it cannot be observed by the requester. We assume that the requester itself always makes effort in the task (i.e., $e_0 = 1$).

Reward payment. The requester pays a reward r_i to each worker i for its contributed data. Given the lack of the ground truth x , this is achieved by the peer prediction mechanism (see, e.g., [Dasgupta and Ghosh, 2013; Liu and Chen, 2016b; 2017; Jin *et al.*, 2017b]) which compares the reported data d_i with the data d_j from a reference worker: a reward b is paid to worker i if and only if $d_i = d_j$:

$$r_i \triangleq b \mathbf{1}_{d_i=d_j}.$$

Then the probability that worker i is rewarded is given by

$$p_i p_j + (1 - p_i)(1 - p_j).$$

Workers' payoff. Based on the crowdsourcing system described above, each worker i 's payoff u_i is the reward r_i paid by the requester minus its *cost* in the task, given by,

$$u_i \triangleq r_i - e_i c.$$

Here the cost c represents how much resource is consumed by worker i (e.g., how much time is spent by worker i) if it makes effort $e_i = 1$ in the task. If worker i make no effort $e_i = 0$, it incurs no cost. Note that the relative weight of the cost c with respect to the reward r_i can be captured by c . We assume that workers have the same cost c which is known to the requester. This assumption is reasonable when the cost c is determined by a uniform market price for working on a task.

Requester's payoff. After collecting all the data \mathbf{d} reported by the workers, the requester aggregates the data \mathbf{d} by making an estimate x_0 of the interested variable X based on \mathbf{d} . Then the utility of crowdsourcing is represented by the correct probability p_c of the estimate x_0 . The requester's payoff u_0 is the crowdsourcing utility (i.e., the correct probability p_c) minus the total reward paid to the workers, i.e.,

$$u_0 \triangleq p_c - \sum_{i \in \mathcal{N}} r_i.$$

4 One-Stage Effort Elicitation

In this section, we study one-stage effort elicitation where the requester collects data once from workers for a task. This serves as basis for sequential multi-stage effort elicitation in the next section.

In this paper, we focus on threshold strategies represented by the threshold p^* : worker i makes effort if and only if its quality p_i is greater than the threshold. Intuitively, it suffices to consider threshold strategies only, since given any other workers' strategies, a worker's payoff is non-decreasing with its quality. To capture the strategic interactions among workers, we consider a Bayesian game where each worker i decides whether to make effort depending on its own quality p_i . We are interested in a Bayesian Nash equilibrium defined as below.

Definition 1 A threshold strategy profile $\{e_i^*(p_i)\}$ is a BNE if

$$E[u_i(e_i^*(p_i)), e_{-i}^*(\mathbf{p}_i)] \geq E[u_i(e_i(p_i)), e_{-i}^*(\mathbf{p}_i)], \forall e_i(p_i).$$

We first consider using a "non-strategic" reference worker in peer prediction. A non-strategic worker is one who always makes effort regardless of other workers' behavior. In this paper, for ease of analysis, we assume that the requester is used as a non-strategic reference worker in peer prediction (except for Proposition 2). The case of using strategic reference workers turns out to be much more complex to analyze than using non-strategic reference workers. This is because the data quality of a reference worker depends on its strategic behavior, which further depends on the behavior of other workers based on their quality.

Proposition 1 The threshold strategy with $p^* = \frac{c}{b(2p_0-1)} + \underline{p}$ is optimal for all workers.

Proof. The reward of worker i when making effort is given by

$$r_i(e_i = 1) = [p_i p_0 + (1 - p_i)(1 - p_0)]b.$$

Thus, the optimal action is to make effort if and only if

$$u_i(e_i = 1) - u_i(e_i = 0) = (p_i - \underline{p})(2p_0 - 1)b - c \geq 0. \quad \square$$

Next we consider using strategic reference workers in peer prediction.

Proposition 2 When the reference workers are strategic, if $F(p)$ is convex, there exist at most two threshold strategies p^* that is a symmetric BNE for all workers.

Proof. The payoff gain of worker i by making effort is given by

$$u_i(e_i = 1) - u_i(e_i = 0) = (p_i - \underline{p})[2E[p_j] - 1] - c.$$

Thus, the threshold p^* is the solution to the equation below:

$$\Delta \triangleq (p^* - \underline{p}) \left[2 \int_{p^*}^{\bar{p}} f(p) p dp + 2F(p^*) \underline{p} - 1 \right] - c = 0.$$

We observe that

$$\frac{\partial \Delta^2}{\partial^2 p^*} = 4f(p^*)(p_0 - p^*) + 2(p^* - \underline{p})[f'(p^*)(p_0 - p^*) - f(p^*)] < 0$$

when $f'(p^*) > 0$. Since Δ is convex, there are at most two solutions to the equation above. \square

5 Sequential Effort Elicitation for Online Quality Learning

In this section, we consider the situation where there are different tasks for which each worker has the same quality. This can capture applications where the tasks are similar in nature (e.g., tagging the same object in different photos). In this situation, it is beneficial for the requester to learn workers' quality from their data, so that it can utilize the learned quality to better perform task allocation and data aggregation. Moreover, the requester can estimate the accuracy of the aggregated data based on the learned quality of workers, which can be useful information. Note that this setting is in contrast to the one-stage setting in the previous section where the requester is oblivious to workers' quality. It is also different from the situations where the requester can elicit workers' quality [Jin *et al.*, 2015; Gong and Shroff, 2017; 2018], since the learned quality is not accurate due to the limited number of data from the workers that can be observed by the requester.

5.1 Online quality learning with sequential peer prediction

We consider a sequential multi-task setting where a new task arrives at each time step. We propose an online quality learning algorithm with sequential peer prediction (as described in Algorithm 1), based on that proposed in [Liu and Liu, 2015]. Specifically, in each time step, it performs one of two functions: exploration or exploitation. In each exploration step, we first estimate the ground truth as $X'(t)$, and then estimate workers' quality $p'_i(t)$ based on the estimated ground truth $X'(t)$, while we reward each worker by b_1 via peer prediction. In each exploitation step, we allocate the task to the optimal worker i^* that has the highest quality⁴, based on the estimated user quality in the past exploration steps, and we reward only the optimal worker by b_2 via peer prediction. Note that compared to most existing online learning problems (such as MAB problems), a major difference of the problem here is that the ground truth is unknown. This challenge is addressed by estimating the ground truth before estimating the workers' quality [Liu and Liu, 2015].

A central design issue for online learning algorithms is to determine when to perform exploration or exploitation, in order to achieve a desired balance between the longterm payoff from exploration and the short-term payoff from exploitation. This tradeoff is quantified in the algorithm by a threshold $h(t)$ given by

$$h(t) = \beta \log t,$$

where

$$\beta = \max\left\{\frac{1}{\epsilon - \alpha}, \frac{1}{\underline{p} - \alpha}\right\}$$

and

$$\epsilon \triangleq \frac{\min_{i \neq j} |p_i - p_j|}{2}$$

and α is any number with $\alpha < \min\{\epsilon, \underline{p}\}$.

⁴In future work, we will generalize the result to allocating the task to the optimal set of workers for minimizing the correct probability p_c of the estimate x_0 .

Algorithm 1: Online quality learning with sequential peer prediction

```

1 Initialize  $t = 0$ ;  $\epsilon(t)$  is the set of exploration time steps up to time  $t$ ;
2 foreach  $t > 0$  do
3   if  $\epsilon(t) > h(t)$  then
4     exploration: allocate task 1 (the task that arrives at  $t = 1$ ) to all workers, and collect data  $D_i(1)$  from the workers;
5     reward each worker  $i$  by peer prediction
6      $r_i = b_1 \mathbf{1}_{D_i(1)=D_0(1)}$ ;
7     estimate the ground truth  $X'(t) = \frac{\sum_{i=1}^N D_i(t)}{N}$ ;
8     estimate each worker' quality
9      $p'_i(t) = \frac{\sum_{i \in \epsilon(t)} \mathbf{1}_{D_i(t)=X'(t)}}{|\epsilon(t)|}$ ;
10    else if  $\epsilon(t) \leq h(t)$  then
11      exploitation: allocate task  $t$  to the worker  $k$  with the highest quality based on the estimated quality  $p'_k(t) = \max_i p'_i(t)$ ;
12      reward the worker  $k$  by peer prediction
13       $r_k = b_2 \mathbf{1}_{D_k(t)=D_0(t)}$ ;
14    end
15  end

```

5.2 Analysis of offline optimal allocation

To provide useful insights, we first study the offline optimal allocation, where the requester knows workers' quality and thus can find the optimal worker. To facilitate further analysis, we first focus on a single exploitation step. As preliminary results, in this paper we assume that workers' actions are static over time (i.e., $e_i(t_1) = e_i(t_2), \forall t_1 \neq t_2$), and we will study dynamic actions of workers in future work.

Proposition 3 *For the offline optimal allocation in an exploitation step, there exists a unique threshold strategy p^* that is a symmetric BNE.*

Proof. When worker i makes effort, its payoff is given by

$$\hat{u}_i(e_i = 1) = F^{N-1}(p_i)(2p_i p_0 - p_i - p_0 + 1)b_2 - c.$$

Otherwise, when worker i makes no effort, its payoff is $\hat{u}_i(e_i = 0) = 0$, since its quality is $p_i(e_i = 0) = \underline{p}$ and thus will not be selected in exploitation steps. Thus the unique threshold p^* is the solution to the equation below:

$$F^{N-1}(p^*)(2p^* p_0 - p^* - p_0 + 1)b_2 - c = 0.$$

We can see that the left-hand side (LHS) above is an increasing function of p^* . Thus there is a unique solution to the equation. \square

We can see that the threshold p^* decreases when the distribution $F(p)$ improves. This is because when other workers are more likely to have high quality, worker i is less likely to be selected as the optimal worker to receive reward in the exploitation step, and thus is less likely to make effort.

Next we turn to offline optimal allocation including both exploration and exploitation steps.

Proposition 4 *For the offline optimal allocation, there exists a unique threshold strategy p^* that is a symmetric BNE.*

Proof. When worker i changes its action from not making effort to making effort, the change of its payoff $\bar{u}_i(e_i = 1) - \bar{u}_i(e_i = 0)$ is given by:

$$\begin{aligned} & \frac{h(t)}{t}(p_i - \underline{p})(2p_0 - 1)b_1 \\ & + (1 - \frac{h(t)}{t})F^{N-1}(p_i)(2p_i p_0 - p_i - p_0 + 1)b_2 - c. \end{aligned}$$

Thus the unique threshold p^* is the solution to the equation below:

$$\begin{aligned} & \frac{h(t)}{t}(p^* - \underline{p})(2p_0 - 1)b_1 \\ & + (1 - \frac{h(t)}{t})F^{N-1}(p^*)(2p^* p_0 - p^* - p_0 + 1)b_2 - c = 0. \end{aligned}$$

We can see that the LHS above is an increasing function of p^* . Thus there is a unique solution to the equation. \square

We can see that the threshold for both exploration and exploitation steps is less than that for exploration steps only or for exploitations steps only. This is because worker i is likely to receive more reward in both exploration and exploitation steps by making effort.

5.3 Analysis of online learning based allocation

Based on the analysis for the offline optimal allocation, we now study the online learning based allocation.

Proposition 5 *The unique threshold strategy that is a symmetric BNE for the offline optimal allocation is a symmetric 2Δ -approximate BNE for the online learning based allocation, where*

$$\Delta \triangleq \sum_{t=1}^T N p_0 \left(\frac{1}{t^2} + \frac{1}{\alpha t^2} \right).$$

Proof. We sketch the proof here and provide the detailed proof in [Gong, 2018]. At each exploration step, each worker's payoff is the same as that in the offline optimal allocation. However, at each exploitation step, each worker's payoff is different from the offline setting, as the optimal worker may not be selected while some non-optimal worker is selected instead. This error in selecting the optimal worker is due to the error in the requester's estimated quality of workers compared to their actual quality. To characterize the gap of a worker's payoff from the offline setting, we first quantify the probability of not selecting the optimal worker, or selecting a particular non-optimal worker. This probability is bounded by the probability that the ordering of workers' estimated quality is wrong, which is bounded as below:

Lemma 1 *The probability $Pr(i^* \neq k)$ that the ordering of workers' estimated quality is different from that of the actual quality in the exploitation step at time t is bounded by*

$$Pr(i^* \neq k) \leq N \left(\frac{2}{t^2} + \frac{1}{\alpha t^2} \right).$$

Based on Lemma 1, the worker's payoff in the online setting is different from that in the offline setting by at most Δ . In particular, we have

$$\bar{u}_i(e_i^*(p_i) = 1) - \Delta \leq u_i(e_i^*(p_i) = 1) \leq \bar{u}_i(e_i^*(p_i) = 1)$$

and

$$\bar{u}_i(e_i^*(p_i) = 0) \leq u_i(e_i^*(p_i) = 0) \leq \bar{u}_i(e_i^*(p_i) = 0) + \Delta.$$

Therefore, if worker i chooses action $e_i^* = 0$ based on the fact $\bar{u}_i(e_i^*(p_i) = 0) > \bar{u}_i(e_i^*(p_i) = 1)$, then we have

$$\begin{aligned} & u_i(e_i^*(p_i) = 0) - u_i(e_i^*(p_i) = 1) \\ & \leq \bar{u}_i(e_i^*(p_i) = 0) + \Delta - (\bar{u}_i(e_i^*(p_i) = 1) - \Delta) < 2\Delta. \end{aligned}$$

If worker i chooses action $e_i^* = 1$ based on the fact $\bar{u}_i(e_i^*(p_i) = 1) > \bar{u}_i(e_i^*(p_i) = 0)$, then we have

$$\begin{aligned} & u_i(e_i^*(p_i) = 1) - u_i(e_i^*(p_i) = 0) \\ & \leq \bar{u}_i(e_i^*(p_i) = 1) - \bar{u}_i(e_i^*(p_i) = 0) > 0. \end{aligned}$$

\square

The gap Δ of a worker's payoff between the online setting and offline setting can be viewed as the worker's "regret" of the requester not using the offline optimal allocation but using the online learning based allocation.

It can be easily seen that $\Delta = \sum_{t=1}^T \left(\frac{1}{t^2} + \frac{1}{\alpha t^2} \right) < \infty$. Therefore, the worker's "regret" is finite, so that the average regret goes to 0 as $T \rightarrow \infty$. This is because the online quality learning is asymptotically optimal, so that the probability of not selecting the optimal worker goes to 0 asymptotically as $T \rightarrow \infty$.

6 Conclusion and Future Work

As data quality is a key metric of data crowdsourcing, it is important to take into account the fact that workers participating in crowdsourcing typically have diverse quality. In this paper, we have studied effort elicitation in data crowdsourcing where workers have diverse quality and know their quality, while it is unknown to the requester. We have shown that for one-stage effort elicitation, a symmetric quality-based threshold strategy is an equilibrium for the workers. We then have proposed an online quality learning mechanism with sequential effort elicitation, for which we show that a symmetric quality-based threshold strategy is an approximate equilibrium.

Based on the preliminary results in this paper, several extensions and generalizations of the problem here will be studied in future work. First, we will study the optimal reward for maximizing the requester's payoff, including that in both exploration and exploitation steps for the online quality learning. Second, we will analyze the equilibrium in the online quality learning when a worker can make different effort over time. Third, we will consider the setting where workers do not know their own quality.

References

[Dasgupta and Ghosh, 2013] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *International World Wide Web Conference (WWW)*, 2013.

- [Duan *et al.*, 2012] Lingjie Duan, Takeshi Kubo, Kohei Sugiyama, Jianwei Huang, Teruyuki Hasegawa, and Jean Walrand. Incentive mechanisms for smartphone collaboration in data acquisition and distributed computing. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2012.
- [Feng *et al.*, 2014] Zhenni Feng, Yanmin Zhu, Qian Zhang, Lionel M Ni, and Athanasios V Vasilakos. Trac: Truthful auction for location-aware collaborative sensing in mobile crowdsourcing. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2014.
- [Gong and Shroff, 2017] Xiaowen Gong and Ness Shroff. Truthful mobile crowdsensing for strategic users with private qualities. In *International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2017.
- [Gong and Shroff, 2018] Xiaowen Gong and Ness Shroff. Incentivizing truthful data quality for quality-aware mobile data crowdsourcing. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2018.
- [Gong, 2018] Xiaowen Gong. Incentivizing quality-based data crowdsourcing. In <https://www.dropbox.com/s/t5vkttxvp28f7f/CS-quality-18-TR.pdf?dl=0>, 2018.
- [Jin *et al.*, 2015] Haiming Jin, Lu Su, Danyang Chen, Klara Nahrstedt, and Jinhui Xu. Quality of information aware incentive mechanisms for mobile crowd sensing systems. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2015.
- [Jin *et al.*, 2016] Haiming Jin, Lu Su, Houping Xiao, and Klara Nahrstedt. INCEPTION: Incentivizing privacy-preserving data aggregation for mobile crowd sensing systems. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2016.
- [Jin *et al.*, 2017a] Haiming Jin, Lu Su, and Klara Nahrstedt. CENTURION: Incentivizing multi-requester mobile crowd sensing. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2017.
- [Jin *et al.*, 2017b] Haiming Jin, Lu Su, and Klara Nahrstedt. Theseus: Incentivizing truth discovery in mobile crowd sensing systems. In *ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2017.
- [Karger *et al.*, 2011] David R Karger, Sewoong Oh, and Devavrat Shah. Budget-optimal crowdsourcing using low-rank matrix approximations. In *IEEE Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2011.
- [Koutsopoulos, 2013] Iordanis Koutsopoulos. Optimal incentive-driven design of participatory sensing systems. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2013.
- [Lee *et al.*, 2015] Donghyeon Lee, Joonyoung Kim, Hyunmin Lee, and Kyomin Jung. Reliable multiple-choice iterative algorithm for crowdsourcing systems. In *ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2015.
- [Liu and Chen, 2016a] Yang Liu and Yiling Chen. A bandit framework for strategic regression. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1821–1829, 2016.
- [Liu and Chen, 2016b] Yang Liu and Yiling Chen. Learning to incentivize: Eliciting effort via output agreement. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2016.
- [Liu and Chen, 2017] Yang Liu and Yiling Chen. Sequential peer prediction: Learning to elicit effort using posted prices. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [Liu and Liu, 2015] Yang Liu and Mingyan Liu. An online learning approach to improving the quality of crowdsourcing. In *ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2015.
- [Luo *et al.*, 2015] Yuan Luo, Nihar B. Shah, Jianwei Huang, and Jean Walrand. Parametric prediction from parametric agents. In *The 10th Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, 2015.
- [Pu *et al.*, 2016] Lingjun Pu, Xu Chen, Jingdong Xu, and Xiaoming Fu. Crowdlet: Optimal worker recruitment for self-organized mobile crowdsourcing. In *IEEE Conference on Computer Communications (INFOCOM)*, 2016.
- [Shah and Zhou, 2015] Nihar Bhadrish Shah and Denny Zhou. Double or nothing: Multiplicative incentive mechanisms for crowdsourcing. In *Conference on Neural Information Processing Systems (NIPS)*, 2015.
- [Tarable *et al.*, 2015] Alberto Tarable, Alessandro Nordio, Emilio Leonardi, and Marco Ajmone Marsan. The importance of being earnest in crowdsourcing systems. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2015.
- [Wang *et al.*, 2016] Weina Wang, Lei Ying, and Junshan Zhang. The value of privacy: Strategic data subjects, incentive mechanisms and fundamental limits. In *ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2016.
- [Yang *et al.*, 2012] Dejun Yang, Guoliang Xue, Xi Fang, and Jian Tang. Crowdsourcing to smartphones: Incentive mechanism design for mobile phone sensing. In *ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2012.
- [Zhang *et al.*, 2016] Honggang Zhang, Benyuan Liu, Hengky Susanto, Guoliang Xue, and Tong Sun. Incentive mechanism for proximity-based mobile crowd service systems. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2016.