Optimal Hiring Strategy in Auction-Based Crowdsourcing Systems

Hongtao Liu[1], Weiran Shen[1](✉), and Yiheng Shen[2]

[1] Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
{ht6,shenweiran}@ruc.edu.cn

[2] Computer Science, Duke University, Durham, NC, USA
ys341@duke.edu

Abstract. We consider an auction-based crowdsourcing system. A requester is faced with a binary choice question and decides to hire workers to answer the question. The workers can ask prices for answering the question and the requester can choose to hire which workers based on their skills and ask prices. We model the problem as a mechanism design problem and characterize the optimal hiring policy. We show that the problem of computing the accuracy of a given set of workers is #P-hard. However, we prove that choosing at most k workers into committee can achieve at least $1/\lceil n/k \rceil$ of the optimal utility. Finally, we also provide a polynomial algorithm for computing the optimal hiring strategy when the number of workers' skill levels is constant.

Keywords: Optimal strategy · Auction · Crowdsourcing system

1 Introduction

Crowdsourcing is a way of hiring a group of participants to finish a certain task together. In a crowdsourcing task, there is a requester and a set of workers. The requester first posts a task advertisement on a certain crowdsourcing platform. A typical advertisement usually contains a brief description of the task and a reward for finishing the task. Then interested workers participate to finish the task and get paid accordingly.

In recent years, crowdsourcing applications have significantly expanded, covering areas such as data labeling for machine learning tasks, information collection, and human subject studies. The requester typically breaks the task into smaller sub-tasks and offers a uniform price for each sub-task. This outsourcing model has both advantages and disadvantages. It can quickly attract many workers, speeding up task completion, but it may also increase task completion costs.

Quality control is another issue that can arise in such a procedure. The workers may have different skill levels and thus may complete each sub-task with varying qualities. To solve the problem, the requester can send each sub-task to multiple workers and then aggregate their solutions to increase the solution quality. However, workers with higher skill levels often cost more to provide higher-quality solutions, while the reward usually remains the same. This hiring policy could lead to an adverse selection problem and deter workers with higher skill levels.

V. Chau et al. (Eds.): IJTCS-FAW 2025, LNCS 15828, pp. 343–355, 2025.
https://doi.org/10.1007/978-981-96-8312-3_26

An auction-based crowdsourcing system can reduce the adverse selection problem. In such a system, each worker sets an asking price, which is the minimum reward the worker must receive if hired. After collecting the asking prices, the requester decides which workers to hire.

In this paper, we consider such an auction-based crowdsourcing system and study the optimal hiring policy of the requester. Our model combines both mechanism design and solution aggregation. We show that computing the optimal set of workers to hire is #P-hard. The difficulty is directly inherited from the hardness of computing the best accuracy of a set of workers. We also put forward an efficient algorithm for computing the optimal hiring strategy when there are only a constant number of different skill levels.

1.1 Related Works

Our paper is related to the crowdsourcing literature. A popular model for crowdsourcing task assignment market is to find a matching between the workers and the tasks [3,6,7,16]. Most of these works model the problem as a bipartite graph, and assign each task to one worker. Different from them, we need to assign the task to an unknown number of workers and infer the true answer from the workers' reported answers. We refer interested readers to [18] for a more comprehensive survey of the history works on truth inference in crowdsourcing. In their setting, the principal needs to learn the quality of the workers while assigning them tasks.

There is another line of work that combines crowdsourcing and mechanism design. Dominic [4] models the crowdsourcing system as an all-pay auction. In their work, each agent exerts effort into each contest with a certain cost and the agent with the highest effort wins. As a follow-up work, Luo [10] considers a similar model, where each agent contributes a certain amount, and the agent with the most contribution wins and gets all the other contributions as their utility. In this paper, however, multiple agents can obtain positive utility by answering the question. There are other works applying peer prediction techniques in designing mechanisms for crowdsourcing platforms [8,14]. Unlike this paper, they do not consider the incentive issues. More recently, Wang [15] applies the auction design into mobile crowdsourcing systems.

In the voting theory, Nitzan [12] first proposes the optimal method for aggregating a set of independent answers with varying accuracies, known as the optimal Nitzan-Paroush weighted majority rule. Later, Ben-Yashar and Nitzan [1] extend their results by adding the prior differences and the dependency of experts' skills on the state of nature. More recently, Berend [2] provides a constant approximation for the error rate of the aggregation rule. In our work, we mainly focus on the "accuracy gain" (accuracy minus 1/2) of the aggregated answer since there is a penalty if the answer is wrong.

In machine learning, the Nitzan-Paroush rule is also applied as the well-known Naive Bayes approach. There are both empirical studies [5,9,13] and theoretical analyses [5,17] showing that this approach performs well on aggregating answers to a binary problem.

2 Model

A requester (she) is faced with a binary choice question. If she answers the question correctly, she gets a reward of R. If she answers incorrectly, her utility is $-R$. Assume that the requester has no knowledge about the question and can only take a random guess (with a 0.5 probability being correct, giving her an expected utility of 0). However, she can hire skillful workers to answer the question through a crowdsourcing system. Assume that the crowdsourcing system uses an auction-based mechanism so that the requester can decide who to hire to answer the question.

Let $N = \{1, 2, \ldots, n\}$ represent the set of all workers. Suppose that each worker i has a skill level that completely determines the probability a_i of his answer being correct. We assume that crowdsourcing system can estimate a_i for the requester based on the worker's past performances, i.e., a_i is known to the requester. Clearly, we have that $a_i \in [0.5, 1]$ since a random guess could already give an accuracy of 0.5. For simplicity, we assume that the workers will try their best to answer the question and that their answers are independent from one another. Therefore, the best accuracy that the requester can get from the set C of hired workers is:

$$A(C) = \sum_{S \subseteq H} \left(\prod_{i \in S} a_i \prod_{j \in H \setminus S} (1 - a_i) \right) \tag{1}$$

We call Eq. (1) the accuracy function.

Each worker also has a cost $c_i \in \mathbb{R}_+$ for answering the question. Suppose that c_i is a random variable that is drawn from cumulative distribution function $F_i(c_i)$ with a density function $f_i(c_i)$. We consider a Bayesian environment and assume that c_i is worker i's private information, whereas the distribution function $F_i(c_i)$ is common knowledge. The workers are asked to report their costs so that the requester can decide who to hire.

Denote by $H = (x, p)$ the requester's hiring policy, where:

- $x : \mathbb{R}_+^n \mapsto \{0, 1\}^n$ is the hiring function that maps the reported costs to a hiring decision, where $x_i = 1$ if and only if worker i is hired;
- $p : \mathbb{R}_+^n \mapsto \mathbb{R}_+^n$ is the payment function that determines the amount of money that the requester pays to each worker.

Worker i's utility u_i is defined as the money he is paid for participation less his cost for answering the question, i.e.,

$$u_i = p_i - c_i x_i, \tag{2}$$

Denote the committee with candidates selected in vector x as $C_x = \{i \in N : x_i = 1\}$. The requester's utility is the expected reward she gets from aggregating the hired workers' answers minus the payments she distributed to the workers:

$$u_0 = R \cdot A(C_x) - R(1 - A(C_x)) - \sum_{i \in N} p_i \tag{3}$$

$$= 2R \cdot A(C_x) - R - \sum_{i \in N} p_i. \tag{4}$$

The complete crowdsourcing procedure is as follows:

1. The requester announces the hiring policy $H = (x, p)$;
2. The workers report their costs c' to the requester (note that it is possible that $c' \neq c$ if misreporting benefits a worker);
3. The requester decides who to hire and the payments by computing $x(c')$ and $p(c')$, and sends the question to the hired workers;
4. Each worker submits his answer to the requester, and gets paid $p_i(c')$.

Following game-theoretic conventions, we use $a = (a_1, a_2, \ldots, a_n)$ to denote the accuracy profile of all workers and $a_{-i} = (a_1, a_2, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ to denote the accuracy profile of all workers except i. Similarly, we can also define c and c_{-i}, respectively.

The workers' goal is to maximize their utilities. Thus they will choose not to participate in the procedure if their utilities are less than 0. Therefore, the requester needs to ensure that the workers' utilities are non-negative if they report truthfully.

Definition 1 (Individual Rationality). *A hiring policy $H = (x, p)$ is individually rational, if $\forall c, i$,*

$$u_i = p_i(c) - c_i x_i(c) \geq 0 \tag{5}$$

And to incentivize the workers to report truthfully, the hiring policy should also satisfy the following incentive compatibility constraint.

Definition 2 (Incentive Compatibility (Truthfulness)). *A hiring policy is said to be incentive-compatible (or truthful), if reporting the cost truthfully is an optimal strategy for each worker. That is, $\forall c_i, c_i', \forall c_{-i}$ and $\forall i$,*

$$p_i(c) - c_i x_i(c) \geq p_i(c_i', c_{-i}) - c_i x_i(c_i', c_{-i}) \tag{6}$$

Thanks to the celebrated revelation principle [11], we can, without loss of generality, focus on the set of truthful hiring policies.

Theorem 1 (Revelation Principle [11]). *For any hiring policy, there exists a truthful hiring policy that implements the same hiring and payment function.*

3 Problem Analysis

In this section, we derive the optimal hiring policy and show the hardness of computing the optimal policy. Note that our setting is very similar to a reverse auction or procurement setting, where a buyer wants to buy an item from multiple sellers. Before we present the optimal hiring policy, we define the following virtual cost function, which is commonly used in the reverse auction design literature.

Definition 3 (Virtual Cost Function). *Suppose worker i's cost is drawn from cumulative distribution function $F_i(c_i)$ with density function $f_i(c_i)$. Then the virtual cost function of worker i is:*

$$\varphi_i(c_i) = c_i + \frac{F_i(c_i)}{f_i(c_i)} \tag{7}$$

A virtual cost function is called regular *if it is monotone increasing.*

Following the standard Myersonian approach, we have the following two results:

Lemma 1. *A hiring policy $H = (x, p)$ is incentive compatible, if and only if the following two conditions hold:*

- *The hiring function $x_i(c_i, c_{-i})$ is monotone decreasing in c_i for all c_{-i};*
- *The payment function $p_i(c_i, c_{-i})$ satisfies:*

$$p_i(c_i, c_{-i}) = \int_0^{c_i} s \, \mathrm{d}x_i(s, c_{-i}), \tag{8}$$

where we fix c_{-i} and view the hiring function $x_i(s, c_{-i})$ as a function of s.

Lemma 2. *For any truthful hiring policy, the requester's utility can be written as:*

$$u_0 = 2RA(C_x) - R - \sum_{i \in N} x_i \varphi_i(c_i). \tag{9}$$

We omit the proofs for the above two lemmas here, as they can be easily derived using standard mechanism design techniques. To optimize the requester's utility, we need to choose a monotone decreasing hiring function x to maximize Eq. (9). The last term in Eq. (9) is easy to maximize if for each i, $\varphi_i(c_i)$ is regular. Even if $\varphi_i(c_i)$ is irregular, we can still apply the so-called "ironing" trick [11] to obtain a monotone increasing version. Define the quantile of a cost c_i as $q_i(c_i) = 1 - F_i(c_i)$ and a function $H_i(q) = (1 - q) \cdot c_i(q)$, where $c_i(q)$ is the cost at quantile q. The ironed virtual cost is the following:

Definition 4 (Ironed Virtual Cost). *The ironed virtual cost of an agent i is defined as:*

$$\overline{\varphi}_i(c_i) = G_i'(q_i(c_i)),$$

where $G_i'(q)$ is the first order derivative (with respect to q) of the convex hull of function H_i:

$$G_i(q) = \max_{\substack{\omega \in [0,1] \\ \alpha \cdot \omega + \beta \cdot (1-\omega) = q}} \{\omega H_i(\alpha) + (1 - \omega) H_i(\beta)\}.$$

Since G_i is convex, the ironed virtual cost is monotone decreasing. It can also be proved that Eq. (9) still holds if we substitute φ_i with $\overline{\varphi}_i$. Based on the previous two lemmas, we present our main mechanism as follows:

Mechanism 1 (Mechanism for Hiring Workers).

1. *Collect the reported cost c_i from all workers.*
2. *Compute each worker's ironed virtual cost $\overline{\varphi}_i(c_i)$.*
3. *Select x that maximizes*

$$u_0 = 2RA(C_x) - \sum_{i \in N} x_i \overline{\varphi}_i(c_i).$$

4. *Set the payment rule as:*

$$p_i(c_i, c_{-i}) = \int_0^{c_i} s \, \mathrm{d}x_i(s, c_{-i}).$$

Theorem 2. *Mechanism 1 is incentive-compatible, individually rational, and maximizes the requestor's utility.*

Proof. The proof directly follows from Lemma 1 and Lemma 2. □

The major challenge for implementing the mechanism is in maximizing the first term in Eq. (9), or equivalently, maximizing the accuracy function. We present the difficult results starting from some simple cases showing that the function is not submodular.

3.1 Simple Case for $n \leq 2$

Let's first consider the easy case with $n \leq 2$. In this case, optimizing Eq. (9) is easy since we can simply enumerate all possible hiring outcomes. There are only 2 outcomes ($x = 0$ and $x = 1$) for $n = 1$, and 4 outcomes ($x \in \{(0,0),(0,1),(1,0),(1,1)\}$) for $n = 2$. A closer investigation into this case, however, shows that the optimal hiring policy is always to hire at most a single worker.

Observation 1. *When $n \leq 2$, the optimal hiring strategy always hires at most 1 worker.*

Proof. The statement is true for $n = 1$. For $n = 2$, we claim that hiring a single worker can achieve the same accuracy as hiring both workers. This option is cheaper and therefore a better strategy. To prove the claim, simply note that the requester will choose to use the answer provided by the worker with a higher accuracy, regardless of whether the two workers' answers agree with each other. Thus, the aggregated answer is always the same as the better worker.

Remark. Since $\ln\left(\frac{0.5+x}{0.5-x}\right)$ is a convex function, similar proof can be extended to show that, if a committee C contains a "dominating high-skill" worker with accuracy a_{big} and other k workers with accuracy $a_1, a_2, \ldots, a_k$ satisfying

$$a_{big} - 0.5 \geq \sum_{i=1}^{k} (a_i - 0.5),$$

we have $A(C) = a_{big}$. Moreover, Observation 1 can be generalized to the following result:

Lemma 3. *When all workers have the same accuracy, i.e., $a_i = a, \forall i$, for any non-negative integer k, hiring $2k+1$ workers yields the same overall accuracy as that of hiring $2k+2$ workers.*

The proof of the Lemma 3 was also deferred to the full version.

We can directly obtain non-submodularity of the accuracy function from Lemma 3.

Corollary 1. *The function A is not submodular.*

4 Hardness of Computing Accuracy

Despite the fact that the accuracy function is not submodular, we try to characterize it in this section since it help us select the workers. When the problem is turned to $n \geq 3$ with different accuracies, the first natural question that comes to our mind is: given the answers from all the workers, how do we optimally aggregate them? Fortunately, the answer to this question is known. We present the result here for completeness.

Lemma 4 (Optimal Weighted Voting Scheme). *The best strategy to aggregate hired workers' answers is equivalent to carrying out weighted voting among the workers with the following voting weight scheme:*

$$w_i \propto \ln \frac{a_i}{1 - a_i}.$$

Proof. See Theorem 1 in [12]. □

The lemma above leads to our next result. Given a binary vector x, it is already a difficult problem to compute the accuracy function $A(C_x)$. Optimizing it is even more challenging.

Theorem 3. *Take each worker's accuracy a_i in a committee C as input, computing $A(C)$ is #P-hard.*

To prove Theorem 3, we reduce the following subset sum problem, which is known to be #P-hard, to our problem.

Lemma 5 (#Subset Sum is #P-hard). *Given a non-negative integer set $I = \{y_1, y_2, \cdots, y_n\}$ and a constant boundary $B = \frac{1}{2}\sum_{i=1}^{n} y_i$, computing $|\{I' \subseteq I \mid \sum_{e \in I'} e > B\}|$ is #P-hard.*

Before proving Theorem 3, we first consider the following lemma:

Lemma 6 (Random Variable Correspondance for Workers). *Suppose in the set C of m hired workers, each worker has accuracies $a_{k_1}, a_{k_2}, \cdots, a_{k_m}$ respectively. Then we can find m random variables, denoted by $X_1, X_2, \cdots, X_m$ satisfying the following:*

$$\Pr\left[X_i = -\ln \frac{a_{k_i}}{1 - a_{k_i}}\right] = 1 - a_{k_i};$$
$$\Pr\left[X_i = \ln \frac{a_{k_i}}{1 - a_{k_i}}\right] = a_{k_i}.$$

If we denote the sum of these variables by $S = \sum_{i=1}^{m} X_i$, we have

$$A(C) = \Pr[S > 0] + \frac{1}{2}\Pr[S = 0].$$

The proof of Lemma 6 was deferred to the full version. Now we are ready to prove Theorem 3.

Proof for Theorem 3. We reduce the #Subset Sum problem to our problem of computing the accuracy. For each instance of #Subset Sum with $B = (y_1 + y_2 + \cdots + y_m)/2$, we assume the answer to the #Subset Sum is K. We construct the following instance for a group of selected workers C.

$$C = \{i_1, i_2, \cdots, i_m\},$$
$$a_{i_j} = 0.5\left[1 + \left(\frac{e^{y_j/2z} - 1}{e^{y_j/2z} + 1}\right)\right], \forall j \in [m]$$
$$z = 1 + \frac{\sum_{j \in [m]} y_j}{\ln\left(1 + \frac{1}{4\, m \cdot 2^m}\right)}.$$

By construction, we have $\ln\left(\frac{a_{i_j}}{1 - a_{i_j}}\right) = \frac{y_j}{2}$. Thus the final accuracy $A(C)$ of the committee C is:

$$\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k})\right) +$$
$$\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u = \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k})\right) \cdot \frac{1}{2}.$$

We can divide the previous formula into two parts, the first part is the probability that the sum of all the workers' corresponding random variables is larger than 0; the second part is half of the probability that the sum of all the workers' corresponding random variables is 0. Our goal is to only keep the first part (otherwise the reduction will be affected), so we consider adding a new worker i_{m+1} to form a new committee. We let the accuracy of the new worker be the following:

$$a_{i_{m+1}} = 0.5\left[1 + \left(\frac{e^{\delta/2z} - 1}{e^{\delta/2z} + 1}\right)\right], 0 < \delta \ll 1.$$

Since $\ln\left(\frac{a_{i_{m+1}}}{1 - a_{i_{m+1}}}\right) = \frac{\delta}{2}$, the random variable corresponding to the new worker X_{m+1} is δ with probability $a_{i_{m+1}}$, and $-\delta$ with probability $1 - a_{i_{m+1}}$. The realized value of all other random variables are all integers. Therefore, when the outcome of X_{m+1} is $-\delta$, the final sum can only be positive if the number of positive variables is at least one more than that of the negative variables. So the final accuracy $A(C \cup \{i_{m+1}\})$ would turn into

$$\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k}) \right) +$$

$$a_{i_{m+1}} \cdot \sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u = \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k}) \right).$$

Thus we have

$$\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k}) \right) = \frac{a_{i_{m+1}} \cdot a_f(C) - \frac{A(C \cup \{i_{m+1}\})}{2}}{a_{i_{m+1}} - \frac{1}{2}}.$$

By letting $\epsilon_j = 0.5 \cdot \left(\frac{e^{y_j/2z} - 1}{e^{y_j/2z} + 1} \right)$, $\epsilon_{\max} = \max_j \{\epsilon_j\}$. We have the following inequality:

$$\begin{aligned}
&\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k}) \right) \\
&= \sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > B}} \left(\prod_{i_j \in C'} (0.5 + \epsilon_j) \cdot \prod_{i_k \notin C'} (0.5 - \epsilon_k) \right) \\
&\le K \cdot \prod_{i_j \in C'} (0.5 + \epsilon_j) \cdot \prod_{i_k \notin C'} (0.5 + \epsilon_k) \\
&\le K \cdot (0.5 + \epsilon_{\max})^m \le \frac{K}{2^m} \cdot (1 + 4\, m \cdot \epsilon_{\max}).
\end{aligned}$$

Similarly, we can get

$$\sum_{\substack{C' \subset C, \\ \sum_{i_u \in C'} y_u > \sum_{i_v \notin C'} y_v}} \left(\prod_{i_j \in C'} a_{i_j} \cdot \prod_{i_k \notin C'} (1 - a_{i_k}) \right) \ge \frac{K}{2^m} \cdot (1 - 4\, m \cdot \epsilon_{\max}).$$

Therefore, we have $z > \frac{\sum_{j \in [m]} y_j}{\ln\left(1 + \frac{1}{4m \cdot 2^m}\right)}$, suppose $\epsilon_{\max}$ is achieved by the worker i_j, we have

$$\epsilon_{\max} = 0.5 \left(\frac{e^{y_j/2z} - 1}{e^{y_j/2z} + 1} \right) < 0.5 \left(e^{y_j/2z} - 1 \right) \le \frac{1}{8\, m \cdot 2^m}.$$

So the lower and upper bounds are:

$$\frac{K}{2^m} \cdot (1 - 4\, m \cdot \epsilon_{\max}) > \frac{K - \frac{1}{2}}{2^m},$$
$$\frac{K}{2^m} \cdot (1 + 4\, m \cdot \epsilon_{\max}) > \frac{K + \frac{1}{2}}{2^m}.$$

Thus we have

$$K - \frac{1}{2} < 2^m \cdot \frac{a_{i_{m+1}} \cdot A(C) - \frac{A(C \cup \{i_{m+1}\})}{2}}{a_{i_{m+1}} - \frac{1}{2}} < K + \frac{1}{2}.$$

Therefore after we computed $A(C)$ and $A(C \cup \{i_{m+1}\})$, we can get K by finding the integer nearest to

$$2^m \cdot \frac{a_{i_{m+1}} \cdot A(C) - \frac{A(C \cup \{i_{m+1}\})}{2}}{a_{i_{m+1}} - \frac{1}{2}}.$$

Thus we have completed the reduction and proved that calculating the accuracy of committee C is #P hard. □

5 Revenue Approximation with Limited-Size Committee

Since it is difficult to compute the utility, we now focus on the case when we only have limited power to compute the optimal strategy with a limited committee size. We studied the utility that this type of committee can achieve to the optimal utility without the size limit. Before presenting our main approximation results, we provide an observation on the corresponding random variable S (in Lemma 6) of a committee. This observation may help readers gain more insights into the accuracy function A.

Observation 2. *The distribution of S is a discrete distribution whose supports are symmetric points:*

$$\{-s_1, -s_2, \ldots, -s_\ell, 0, +s_\ell, +s_{\ell-1}, +s_{\ell-2}, \ldots, +s_1\},$$

where $s_i > 0, \forall i \in [\ell]$. *Moreover, we have* $\Pr[S = +s_j] > \Pr[S = -s_j], \forall j \in [\ell]$.

The proof of Observation 2 was deferred to full version.

Theorem 4. *The optimal hiring strategy with k workers gives the requester a utility at least* $1/\lceil \frac{n}{k} \rceil \cdot U^*$, *where U^* is the maximum utility with unlimited committee size.*

Before proving this theorem, we first consider a useful lemma:

Lemma 7. *For any two groups of workers C_1 and C_2,* $A(C_1 \cup C_2) - 1/2 \leq A(C_1) + A(C_2) - 1$.

The proof of Lemma 7 was also deferred to the full version. With the above lemma, we are ready to prove Theorem 4.

Proof of Theorem 4. For simplicity, we define the "accuracy gain" of a worker group C as $A(C) - 1/2$. Suppose that the optimal hiring strategy contains n^* workers, we can divide the whole committee into $q = \lceil \frac{n^*}{k} \rceil$ disjoint groups, each group contains at most k workers. Denote these groups by $G_1, G_2, \ldots, G_q$. Suppose that the maximum utility with a committee of size at most k is U'_k. Then the utility of hiring each group is at most U'_k. Then we have

$$\begin{aligned} U^* &= \left[(2A\left(\bigcup_{i=1}^{q} G_i\right) - 1\right] \cdot R - \sum_{i=1}^{q} \bar{\varphi}(G_i) \\ &= R \cdot \sum_{i=1}^{q}(2A(G_i) - 1) - \sum_{i=1}^{q} \bar{\varphi}(G_i) \\ &\leq \sum_{i=1}^{q} (R \cdot (2A(G_i) - 1) - \bar{\varphi}(G_i)) \\ &\leq U'_k \cdot q = U'_k \left\lceil \frac{n^*}{k} \right\rceil \leq U'_k \left\lceil \frac{n}{k} \right\rceil, \end{aligned}$$

where the second equation follows from Lemma 7, and the second inequality is due to the fact that $|G_i| \leq k$. Therefore, we have

$$U'_k \geq \frac{U^*}{\lceil \frac{n}{k} \rceil},$$

completing the proof. □

6 Algorithm for Selecting Workers with L Possible Accuracies

In this section, we consider a specific case where there are only a constant number of possible accuracies and propose an algorithm for selecting the optimal committee. Let L be the number of possible accuracies. We can categorize the workers into L groups with each group containing workers with the same accuracy. Denote the accuracies by $a^1, a^2, \ldots, a^L$, and the number of workers in each group by $n^1, n^2, \ldots, n^L$, we give an algorithm that returns the optimal committee in polynomial time. The algorithm is as follows:

1. Sort workers in each group according to their ironed virtual costs in ascending order. Let $p_{i,j}$ be the sum of the first j workers in group i.
2. Compute the accuracy $A_{w_1, w_2, \ldots, w_L}$ of a committee with w_i workers in each group i.
3. Compute $(w_1^*, w_2^*, \ldots, w_L^*) =$

$$\arg\max_{w_1, \ldots, w_L} \left\{ (2A_{w_1, w_2, \ldots, w_L} - 1) \cdot R - \sum_{i=1}^{L} p_{i, w_i} \right\}.$$

   If the maximum value is positive, select the first w_i workers in group i into the committee. Otherwise, do not select any worker.

Note that in the second step, the w_i workers in group i has the same accuracy. Thus in the third step, we will clearly choose the w_i workers with the lowest virtual costs.

Theorem 5. *The above algorithm runs in polynomial time and induces a truthful mechanism.*

Proof. There are only $\prod_{i=1}^{L}(n^i+1)$ different combinations in step 3. The only missing point is how to compute the $A_{w_1,w_2,\ldots,w_L}$ in the second step for any combination of $\{w_i\}$ in polynomial time.

We construct an L-dimension "matrix" M with $n^1 \times n^2 \times \cdots \times n^L$ elements. Each element $m_{w_1,w_2,\cdots,w_L}$ stores the distribution of the random variable corresponding to (by Lemma 6) the committee with w_i workers in layer i. For any random variable, each group of workers can contribute a value from the following set to the sum:

$$\left\{ -n^i \cdot \ln \frac{a^i}{1-a^i}, -(n^i-1) \cdot \ln \frac{a^i}{1-a^i}, \ldots, (n^i-1) \cdot \ln \frac{a^i}{1-a^i}, n^i \cdot \ln \frac{a^i}{1-al^i} \right\},$$

thus there are at most $\prod_{i=1}^{L}(2n^i+1)$ possible values of the random variable.

We compute the distribution for $m_{w_1,w_2,\cdots,w_L}$ using the distribution for elements with smaller indices. For example, when computing the distribution for $m_{w_1,w_2,\ldots,w_L}$, we choose an arbitrary $i \in [L]$ where $w_i \geq 1$. We begin with the distribution for the element $m' = m_{w_1,w_2,\ldots,w_i-1,\ldots,w_L}$, where the i-th dimension index is decreased by 1. The support of m' is on at most $G = \prod_{i=1}^{L}(2n^i+1)$ values. We enumerate all its support x with probability $p(x)$. We then add two resulting points $x + \ln \frac{a^i}{1-a^i}$ with probability $p(x) \cdot a^i$ and $x - \ln \frac{a^i}{1-a^i}$ with probability $p(x) \cdot (1-a^i)$ to the distribution of $m_{w_1,w_2,\ldots,w_L}$.

Therefore, each distribution takes $O\left(\prod_{i=1}^{L}(2n^i+1)\right)$ computations and there are in total $\prod_{i=1}^{L}(n^i+1)$ accuracies to compute. Therefore, the total running time of the algorithm is $O((2n)^L \cdot n^L) = O(2^L \cdot n^{2L})$, which is polynomial in n since L is a constant.

To show that the induced mechanism is truthful, it suffices to show that the probability of a worker being chosen decreases as their cost increases. For any worker i, the ironed virtual cost function $\bar{\varphi}(c_i)$ is monotone increasing in c_i. So if c_i increases, worker i is more likely to be placed at the front, and all $p_{i,j}$ are weakly increasing. As a result, the committee containing worker i is less likely to be chosen according to step 3 of our algorithm. □

7 Conclusion

In this paper, we consider the mechanism design problem where a platform needs to hire a group of workers to answer a binary question. We characterize truthful mechanisms and give an optimal mechanism. We also show that computing the overall accuracy for a given group of workers is #P-hard, and give approximation results when the committee size is capped by a constant. Finally, when the workers' skill levels fall in constant layers, we give a polynomial algorithm for computing the optimal hiring strategy.

Since our paper only considers a dichotomous choice question for the requestor, a natural extension of our work is to consider a non-binary question with more than 2 answers. It is also interesting to consider the interdependency among the workers' answers and the accuracy of workers are affected by the state of nature (true answer).

References

1. Ben-Yashar, R.C., Nitzan, S.I.: The optimal decision rule for fixed-size committees in dichotomous choice situations: the general result. Int. Econ. Rev. 175–186 (1997)
2. Berend, D., Kontorovich, A.: Consistency of weighted majority votes. In: Advances in Neural Information Processing Systems, vol. 27 (2014)
3. Dickerson, J.P., Sankararaman, K.A., Srinivasan, A., Xu, P.: Assigning tasks to workers based on historical data: online task assignment with two-sided arrivals. In: International Conference on Autonomous Agents and Multiagent Systems (AAMAS) (2018)
4. DiPalantino, D., Vojnovic, M.: Crowdsourcing and all-pay auctions. In: Proceedings of the 10th ACM Conference on Electronic Commerce, pp. 119–128 (2009)
5. Domingos, P., Pazzani, M.: Beyond independence: conditions for the optimality of the simple bayesian classifier. In: Proc. 13th International Conference Machine Learning, pp. 105–112. Citeseer (1996)
6. Goel, G., Nikzad, A., Singla, A.: Mechanism design for crowdsourcing markets with heterogeneous tasks. In: Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, vol. 2 (2014)
7. Hassan, U.U., Curry, E.: A multi-armed bandit approach to online spatial task assignment. In: 2014 IEEE 11th International Conference on Ubiquitous Intelligence and Computing and 2014 IEEE 11th International Conference on Autonomic and Trusted Computing and 2014 IEEE 14th International Conference on Scalable Computing and Communications and Its Associated Workshops, pp. 212–219. IEEE (2014)
8. Kamar, E., Horvitz, E.: Incentives for truthful reporting in crowdsourcing. In: AAMAS, vol. 12, pp. 1329–1330 (2012)
9. Langley, P., Iba, W., Thompson, K., et al.: An analysis of bayesian classifiers. In: AAAI, vol. 90, pp. 223–228. Citeseer (1992)
10. Luo, T., Das, S.K., Tan, H.P., Xia, L.: Incentive mechanism design for crowdsourcing: an all-pay auction approach. ACM Trans. Intell. Syst. Technol. (TIST) **7**(3), 1–26 (2016)
11. Myerson, R.B.: Optimal auction design. Math. Oper. Res. **6**(1), 58–73 (1981)
12. Nitzan, S., Paroush, J.: Optimal decision rules in uncertain dichotomous choice situations. Int. Econ. Rev. 289–297 (1982)
13. Pazzani, M.J.: Searching for dependencies in bayesian classifiers. In: Pre-proceedings of the Fifth International Workshop on Artificial Intelligence and Statistics, pp. 424–429. PMLR (1995)
14. Satzger, B., Psaier, H., Schall, D., Dustdar, S.: Auction-based crowdsourcing supporting skill management. Inf. Syst. **38**(4), 547–560 (2013)
15. Wang, Y., Cai, Z., Zhan, Z.H., Gong, Y.J., Tong, X.: An optimization and auction-based incentive mechanism to maximize social welfare for mobile crowdsourcing. IEEE Trans. Comput. Soc. Syst. **6**(3), 414–429 (2019)
16. Xu, P., Srinivasan, A., Sarpatwar, K.K., Wu, K.L.: Budgeted online assignment in crowdsourcing markets: theory and practice. In: AAMAS, pp. 1763–1765 (2017)
17. Zhang, H.: The optimality of naive bayes. Aa **1**(2), 3 (2004)
18. Zheng, Y., Li, G., Li, Y., Shan, C., Cheng, R.: Truth inference in crowdsourcing: is the problem solved? Proc. VLDB Endowment **10**(5), 541–552 (2017)