Incentivizing Sequential Crowdsourcing Systems

Extended Abstract

Yuan Luo
School of Science and Engineering, The Chinese University of Hong Kong (Shenzhen), Shenzhen, China
Shenzhen Institute of Artificial Intelligence and Robotics for Society, Shenzhen, China
luoyuan@cuhk.edu.cn

ABSTRACT

A crowdsourcing system such as Amazon's Mechanical Turk allows a crowdsource campaign initiator to recruit a large number of workers to accomplish a task. The proper design of such a crowdsourcing system becomes very challenging when the task involves multiple interdependent micro-tasks, and the initiator wants the task to be completed with the minimal cost and a high probability of success. In this paper, we address this challenge by designing an `EI` (Effort Incentivization) mechanism, which utilizes the peer effect to incentivize workers to act according to the initiator's best interest. We prove that `EI` is Bayesian incentive compatible and Bayesian individually rational. Our analysis shows that when there are multiple sequential interdependent micro-tasks, the initiator should provide higher rewards to those workers responsible for completing later stage micro-tasks. When there is a flexibility regarding the worker assignment to each micro-task, the initiator should assign fewer workers to later stage micro-tasks to minimize the initiator's overall payment. Numerical results show that our proposed `EI` mechanism can reduce the initiator's total payment by more than 70%, compared to a fixed reward mechanism. By optimizing the numbers of workers assigned to different interdependent micro-tasks, the initiator can reduce the total payment by up to 50% compared to a random assignment scheme.

KEYWORDS

Crowdsourcing systems; Incentive mechanism; performance guarantees

ACM Reference Format:
Yuan Luo. 2023. Incentivizing Sequential Crowdsourcing Systems: Extended Abstract. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), London, United Kingdom, May 29 – June 2, 2023*, IFAAMAS, 3 pages.

1 INTRODUCTION

Crowdsourcing, which often involves a large number of non-expert agents in completing rather complicated tasks, is becoming an increasingly popular approach to exploit the wisdom of crowd for creating new products and services [11]. In crowdsourcing systems, a crowdsource campaign initiator ("initiator" hereafter) can recruit hundreds of workers through the Internet for accomplishing tasks, and the workers obtain rewards during this process [19]. Typical examples of crowdsourced tasks include image labeling [14] and prediction [7]. For the ease of exposition, we refer to the initiator as "she" and a worker as "he" in this paper.

Due to the heterogeneity of workers in terms of their costs and efforts when completing the tasks, the initiator has an uncertainty regarding the task completion result, which leads to challenges in terms of deciding the proper rewards to workers (which should often be announced to the workers beforehand). Specifically, the rewards should not just cover the workers' costs but also encourage the workers to exert their maximum efforts to accomplish the tasks. However, the workers may have incentives to shirk (not exerting efforts), as it can be very costly for the initiator to effectively monitor each worker's action. This is also called the *hidden action problem* in the mechanism design literature.

To address the shirking problem in crowdsourcing systems, researchers proposed several approaches [3, 5, 9, 12] to incentivize workers to exert the amount of efforts desired by the initiator. The basic idea of these approaches is to reward the worker based on the difference between the actual outcome (*e.g.*, realized stock price after the prediction task is completed) and the task result that the worker produced (e.g., the workers' prediction of the stock price). These approaches rely on the critical assumption that the actual outcome can be verified by the initiator. Another key assumption is that the worker's response and the task outcome can be easily compared. All the above mentioned literature has focused on the incentive issues when workers work on *independent* micro-tasks.

In this paper, we will consider a more complicated yet practical scenario, where the initiator wants workers to complete a sequence of *interdependent* micro-tasks [2, 6, 15, 17]. The initiator divides the workers into several groups and assigns one group to each micro-task. While workers in the same group accomplish the same micro-task simultaneously, workers in different groups complette the micro-tasks sequentially (according to a given order). For example, Bernstein *et al.* in [2] considered a Find-Fix-Verify (FFV) workflow to correct a text. In more detail, the FFV workflow splits a complex text editing task into three interdependent micro-tasks: *Find*, *Fix*, and *Verify*. In the Find stage, a group of workers identify candidate mistakes in the sentences. In the Fix stage, another group of workers correct these mistakes. The final group of workers verify the fixed results in the Verify stage.

The FFV workflow example shows that the performance of the whole task depends on the efforts of workers in all three micro-tasks. Furthermore, the initiator may only observe the performance of the whole task (*e.g.*, how many mistakes have been corrected at the end of the FFV workflow), as checking each micro-task's outcome may be very time and money consuming [4, 10, 16]. Hence, the initiator cannot utilize the outcome of each micro-task, as proposed in [3, 5, 9, 12], to evaluate the workers' individual performances. The key challenge that the initiator needs to solve is how to incentivize workers to perform the interdependent micro-tasks without monitoring the outcomes of individual micro-tasks, so that to obtain

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom.

a high quality overall result with a minimum cost (total payment to workers). In the rest of this paper, we will use “cost" and “total payment to workers" interchangeably.

Several studies in [1, 8, 13, 18] proposed some mechanisms to tackle the sequential hidden action problem in the context of supply chain management, with the common assumption there is only *one* worker working on each micro-task. In the crowdsourcing scenario involving *multiple* workers per micro-task, however, the incentive design problem is more challenging, as a worker who shirks will not incur any cost and may only experience a slight group performance degradation (if most other workers working on the same micro-task exert their efforts) [10].

Again this background, this paper formulates and optimally solves the incentive mechanism design problem in a sequential crowdsourcing system. Specifically, we design an “`EI`" (Effort Incentivization) mechanism that minimizes the initiator’s total reward to workers, while guaranteeing a good performance of the whole task by incentivizing all participating workers to exert effort. In `EI`, the initiator only needs to verify the whole task outcome, instead of checking the outcome of each micro-task. Furthermore, `EI` provides guidance for the initiator to optimally allocate different numbers of workers to finish different interdependent micro-tasks.

2 MODEL AND OBSERVATIONS

We consider a scenario where an initiator splits a complex task into a set $\mathcal{J} = \{1, \dots, J\}$ of J micro-tasks. These micro-tasks are interdependent and need to be accomplished sequentially. The interactions between the initiator and the workers are as follows. The initiator first announces the set of micro-tasks and a payment that is calculated based on `EI` for a worker in $j \in \mathcal{J}$. Then, each worker in micro-task j chooses his effort level (*i.e.*, shirking or not) based on the reward function R_j, all his predecessors’ effort actions, and the assumption that all other workers in the same micro-task will exert effort to accomplish the same micro-task j. After all micro-tasks have been finished by workers, the initiator checks the outcome of the whole task and pays workers based on the announced payment function. As micro-tasks are interdependent and the whole task’s success requires that all micro-tasks have been performed successfully, the whole task’s outcome is a function of all workers’ actions. The probability of whole task success increases with the number of workers exerting efforts. Note that the outcome of whole task has only two states: the whole task failure (*i.e.*, $o^{\mathrm{WT}}(\boldsymbol{a}) = 0$) and the whole task success (*i.e.*, $o^{\mathrm{WT}}(\boldsymbol{a}) = 1$). We present the high-level structure of `EI` in Algorithm 1.

Algorithm 1 Effort Incentivization (`EI`) Mechanism

1: The initiator announces the reward function $R_j : [0,1] \times \mathbb{R}_+^J$ for a worker in $j \in \mathcal{J}$;
2: **for** Micro-task $j = 1$ to J **do**
3: Each worker $n \in \mathcal{N}_j$ in micro-task j chooses his effort level (*i.e.*, shirking or not) based on the reward function R_j, all his predecessors’ effort actions, and the assumption that all other workers in the same micro-task will exert effort to accomplish the same micro-task j;
4: **end for**
5: The initiator verifies the outcome of the whole task and transfers the rewards to the workers accordingly.

Our results give us the following observation:

(a) Under the homogeneous workers and homogeneous interdependent micro-tasks scenario, the reward increases with the worker’s cost c, since a higher c gives the worker more incentive to shirk. The reward increases with the micro-task’s index j. This is because, given this reward function, the shirking of a worker will trigger all later workers to shirk. With a larger micro-task’ index j, the number of workers who make decisions after a worker in micro-task j decreases. Then the impact of the worker’s shirking on the probability of whole task success and the probability of the worker’s receiving reward decreases. Hence, a worker working on a later micro-task has a higher incentive to shirk.

(b) Consider the homogeneous workers and heterogeneous interdependent micro-tasks scenario, with a given set of worker allocation schemes that achieve the same probability of whole task success. In order to minimize the initiator’s total payment, the initiator should select the worker allocation scheme from the set such that the number of allocated workers is non-increasing in the micro-task index.

(c) Consider the heterogeneous workers and heterogeneous interdependent micro-tasks scenario, with a given number of workers per micro-tasks. The initiator should select the worker allocation scheme that assigns the lower costs of workers to later stage micro-tasks and the higher costs of workers to earlier stage micro-tasks. Specifically, the workers are allocated to the sequential micro-tasks based on the descending order of their costs. When the number of micro-tasks is J, we have $\min_{n\in\mathcal{N}_j} c_{j,n} \geq \max_{n\in\mathcal{N}_i} c_{i,n}, \forall i, j \in \mathcal{J}, i > j$, where $\mathcal{N}_j$ denotes the set of the workers assigned to accomplished micro-task $j \in \mathcal{J}$ and $c_{j,n}$ denotes the cost of the worker n in the micro-task $j \in \mathcal{J}$.

3 CONCLUSION

We study the hidden action problem and the workers allocation problem in the sequential crowdsourcing. Our study generalizes two groups of previous studies that look at only one of these problems, respectively. We design an `EI` mechanism to incentivize the workers to exert efforts, so that the initiator can achieve a target probability of whole task success with a minimum payment to the workers. Our analysis shows that the initiator should select the worker allocation scheme such that the number of allocated workers is non-increasing in the micro-task index. Furthermore, the workers with lower cost should be assigned to accomplish later stage micro-tasks.

ACKNOWLEDGMENTS

This work was supported in part by the Shenzhen Institute of Artificial Intelligence, the National Natural Science Foundation of China under Grant No. 62102343 and the Shenzhen Science and Technology Program under Grant No. JCYJ20220818103006012.

REFERENCES

[1] Moshe Babaioff, Michal Feldman, Noam Nisan, and Eyal Winter. 2012. Combinatorial agency. *Journal of Economic Theory* 147, 3 (2012), 999–1034.

[2] Michael S Bernstein, Greg Little, Robert C Miller, Björn Hartmann, Mark S Ackerman, David R Karger, David Crowell, and Katrina Panovich. 2015. Soylent: a word processor with a crowd inside. *ACM Communications* 58, 8 (2015), 85–94.

[3] Ruggiero Cavallo and Shaili Jain. 2012. Efficient crowdsourcing contests. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 677–686.

[4] Yeon-Koo Che and Seung-Weon Yoo. 2001. Optimal incentives for teams. *American Economic Review* (2001), 525–541.

[5] Jenny J Chen, Natala J Menezes, Adam D Bradley, and TA North. 2011. Opportunities for crowdsourcing research on amazon mechanical turk. *Interfaces* 5, 3 (2011).

[6] Yili Fang, Pengpeng Chen, et al. 2022. Hint: harnessing the wisdom of crowds for handling multi-phase tasks. *Neural Computing and Applications* (2022), 1–23.

[7] Jörg Gottschlich and Oliver Hinz. 2014. A decision support system for stock investment recommendations using collective wisdom. *Decision support systems* 59 (2014), 52–62.

[8] Marina Halac, Elliot Lipnowski, and Daniel Rappoport. 2021. Rank uncertainty in organizations. *American Economic Review* 111, 3 (2021), 757–86.

[9] Matthias Hirth, Tobias Hoßfeld, and Phuoc Tran-Gia. 2013. Analyzing costs and accuracy of validation mechanisms for crowdsourcing platforms. *Mathematical and Computer Modelling* 57, 11 (2013), 2918–2932.

[10] Bengt Holmstrom. 1982. Moral hazard in teams. *The Bell Journal of Economics* (1982), 324–340.

[11] Jeff Howe. 2008. *Crowdsourcing: How the power of the crowd is driving the future of business.* Random House.

[12] Yuan Luo, Nihar B Shah, Jianwei Huang, and Jean Walrand. 2018. Parametric prediction from parametric agents. *Operations Research* 66, 2 (2018), 313–326.

[13] Christopher W Miller and Insoon Yang. 2015. Optimal dynamic contracts for a large-scale principal-agent hierarchy: a concavity-preserving approach. *arXiv preprint arXiv:1506.05497* (2015).

[14] Bryan C Russell, Antonio Torralba, Kevin P Murphy, and William T Freeman. 2008. LabelMe: a database and web-based tool for image annotation. *International Journal of Computer Vision* 77, 1-3 (2008), 157–173.

[15] Nasim Sabetpour, Adithya Kulkarni, Sihong Xie, and Qi Li. 2021. Truth discovery in sequence labels from crowds. In *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE, 539–548.

[16] Roland Strausz. 1999. Efficiency in sequential partnerships. *Journal of Economic Theory* 85, 1 (1999), 140–156.

[17] Long Tran-Thanh, Trung Dong Huynh, Avi Rosenfeld, Sarvapali D Ramchurn, and Nicholas R Jennings. 2015. Crowdsourcing complex workflows under budget constraints. (2015).

[18] Eyal Winter. 2010. Transparency and incentives among peers. *The RAND Journal of Economics* 41, 3 (2010), 504–523.

[19] Yu Zhang and Mihaela van der Schaar. 2012. Reputation-based incentive protocols in crowdsourcing applications. In *IEEE International Conference on Computer Communications (INFOCOM)*. 2140–2148.