Multi-Armed Bandit Approach to Task Assignment across Multi Crowdsourcing Platforms

Yunyi Xiao, Yu Yamashita, Hiroyoshi Ito, Masaki Matsubara, Atsuyuki Morishima
University of Tsukuba
s1913024@klis.tsukuba.ac.jp, yu.yamashita.2021b@mlab.info, ito@slis.tsukuba.ac.jp, masaki@slis.tsukuba.ac.jp, morishima-office@ml.cc.tsukuba.ac.jp

Abstract

Many existing optimization approaches deal with task assignments on one particular platform. This paper addresses the problem of task assignments to multiple crowdsourcing platforms. We model this problem as a Multi-Armed Bandit (MAB) Problem for the following reasons. First, the optimal platform is not trivial in general. Second, it can easily control the task assignment policy by changing the setting for the MAB Problem. This paper overviews the approach and reports our preliminary result, which clearly supported that (1) choosing the optimal platform based on prior knowledge is not necessary easy and (2) our approach for multi-platform task assignment problems is promising.

Introduction

Many optimization strategies have been proposed for the assignment of crowdsourcing tasks to workers on a single crowdsourcing platform (Hettiachchi, Kostakos, and Goncalves 2022; Chittilappilly, Chen, and Amer-Yahia 2016; Li et al. 2016). Besides, comparative studies and studies on combining different platforms for complex workflows have made progress. (Peer et al. 2017; Aizawa et al. 2020). Nonetheless, there is little awareness of the significance of optimizing task assignments across multiple platforms.

Generally, the same task leads to different results on different platforms since the effectiveness of platforms for the task differs. Hence, the requester should choose an appropriate platform for their tasks, considering factors such as language and culture.

However, many difficulties exist when predicting the optimal platform neither by using prior knowledge nor by conducting research; in general, there are multiple contexts for task submissions, such as types of tasks, the attribute of preferable workers, and the time they complete the tasks.

The prior knowledge cannot perfectly handle the interactions of various contexts. For example, given two choices for platforms (Amazon Mechanical Turk (AMT) and Yahoo! Japan), the requester may choose Yahoo! Japan for a Chinese dishes labeling task since Japan is closer to China in terms of culture and geography. Nevertheless, our experiment results have shown that choosing a good platform is a non-trivial task; the difference between the effectiveness of platforms can change as task difficulty changes. When the task is easy, workers on AMT and Yahoo! Japan show similar performance (see E_1 in Table 1). Also, conducting research in advance will cause extravagant expenses because the number of combinations of contexts is innumerable.

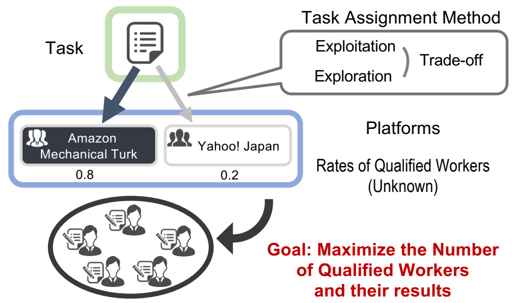

Figure 1: Modeling a task assignment problem on multiple platforms as a Multi-Armed Bandit Problem.

This paper focuses on the optimization problem on multi-crowdsourcing platforms and models it as a Multi-Armed Bandit (MAB) Problem (Figure 1). The modeling gives two advantages. First, it can find the optimal arm (platform) without any prior knowledge or investigations. Second, various algorithms in different settings could establish their policies for choosing platforms, which ensures the satisfaction of different needs. To verify the validity of our approach, we collected data by assigning tasks on crowdsourcing platforms and conducted three experiments using these data. The collected data indicates that the optimal platform is not trivial in general, even with prior knowledge. The experiments supported the effectiveness of our approach.

Modeling the Problem as a MAB Problem

In the Multi-Armed Bandit Problem, we have K arms, each of which is associated with unknown distributions of rewards delivered by the arm. The arms and distributions are represented by environment E as a whole. The gambler can play the arms at T rounds; he plays one arm $i(t)$ per round t and obtains the reward $r_i(t)$, iteratively. The objective is to

Dataset ID	Task	#Tasks	Platform	Commission per Task	#Qualified Workers
E_1 (for Expt 1-MAB)	Labeling Chinese Dishes (Easy)	200	Yahoo! Japan	14.3JPY	160
			AMT	$0.04	144
E_2 (for Expts 2-MAB, 2-BMAB)	Labeling Chinese Dishes (Difficult)	200	Yahoo! Japan	14.3JPY	126
			AMT	$0.04	62

Table 1: Real-world datasets for building experiment environments

maximize the sum of the rewards after all rounds end.

Arm Setting. In our context, each crowdsourcing platform is modeled as an arm. We have tasks to be submitted to one of the platforms. Each task has two sections: (1) qualifying questions for which we know correct answers and (2) other questions for collecting answers.

Round Setting. There are several ways to define a round for pulling an arm. Here are two examples. (1) Each round is represented by the submission of a single task to a particular platform. (2) Each round is represented by a set of tasks completed on a particular platform in a time window.

Reward Setting. The setting of the reward should be determined depending on the goal of the applications. For example, in our preliminary experiment (see below), if a worker gave a correct answer to qualifying questions and completed the whole task within a limited time, we infer that this worker is qualified. Consequently, we receive the reward. If we want to involve the throughput of task completion, the reward setting should consider the factor.

Algorithms. If the requester aims purely to maximize the number of qualified workers, algorithms for the basic MAB setting such as Annealing ϵ-greedy (AE), Thompson Sampling (TS), and Upper Confidence Bound (UCB) will be applied to the problem. If the requester wanted to consider the difference in the task submission cost (e.g., different commissions of platforms), we would adopt Budget-limit Multi-Armed Bandit (BMAB) setting and algorithms like Knapsack–Based Upper Confidence Bound Exploration and Exploitation (KUBE) (Tran-Thanh et al. 2012). Under the BMAB setting, additionally, pulling an arm i will lead to a cost c_i. The total Budget is denoted by B. The total number of rounds T will be decided by the algorithm.

Preliminary Results

We conducted three experiments (1-MAB, 2-MAB, and 2-BMAB) with two real-world datasets E_1 and E_2 that represent two environments. Table 1 shows E_1 and E_2, which we obtained by assigning tasks to workers on two real-world crowdsourcing platforms to label Chinese dishes with different difficulties. We set the incentives to workers based on the minimum wages of the countries and the limited completion time to 15 minutes. For the collection of data, the tasks only contained qualifying questions without losing the generality.

Both environments have $K = 2$ arms (Yahoo! Japan as arm $i = 1$, AMT as arm $i = 2$). In E_1, the two arms have similar reward distributions while there are separated reward distributions in E_2, which shows that which platform is good for a given set of tasks is not trivial. In the experiment, the round is modeled as the submission of a single task to one of the platforms. The reward $r_i(t) = \mathbf{1}(a(t) = a^*)\mathbf{1}(t_c(t) \leq t_c^*)$. The two indicator functions show whether the answer $a(t)$ is the same as the correct answer a^*, and whether the completion time $t_c(t)$ exceeds the limitation t_c^*, respectively. In the basic MAB setting, we applied AE, UCB, and TS with the total number of rounds $T = 200$. In SMAB setting, we applied KUBE with $c_1 = 14$, $c_2 = 5$, $B = 1000$.

Algorithm	Total Rounds	CR (Qualified Worker)	Total Cost	Selection Rate (arm 1 / arm 2)
Random (MAB)	200	98	1891	$99/101$
AE		115	2476	$164/36$
UCB		120	2710	$190/10$
TS		**125**	2692	$\mathbf{188/12}$
Oracle (MAB)		153	2557	$173/27$
Random (BMAB)	110	49	1000	$50/60$
KUBE	169	**56**	998	$\mathbf{17/152}$
Oracle (BMAB)	108	83	999	$51/57$

Table 2: Results of algorithms in E_2. (Top-half part: Expt 2-MAB. Bottom-half part: Expt 2-BMAB)

Results. Table 2 summarizes the results of Expts 2-MAB and 2-BMAB (1-MAB is omitted due to space limitation). We have the following observations. First, all algorithms resulted in larger amounts of cumulative rewards (CR) than random results. Note that the amount of CR using TS still surpassed the random result in Expt 1-MAB in spite of the little differences in reward distributions. Second, there are a few differences in the performances of algorithms in this environment. UCB and TS did not show distinguished differences in their performances but both outperformed AE. Third, the policies of selecting arms established by algorithms reverse under the basic MAB setting and under the BMAB setting. For the MAB setting, algorithms have selected $i = 1$ more and $i = 1$ is considered to be the optimal arm. Conversely, KUBE has selected $i = 2$ more and $i = 2$ became the optimal arm under the BMAB setting. After the same number of rounds, MAB algorithms found more qualified workers while BMAB led to a smaller cost.

This research was approved by IRB of University of Tsukuba.

Future Work

For future work, we will (1) extend this work into a non-stationary environment by applying time-aware MAB, and (2) explore interactions of tasks and algorithms.

Acknowledgment. This work was partially supported by JSPS KAKENHI Grant Number 22H00508 and 21H03552.

Reference

Hettiachchi, D.; Kostakos, V.; and Goncalves, J. 2022. A Survey on Task Assignment in Crowdsourcing. *ACM Comput. Surv.*, 55: 1–35.

Chittilappilly, A. I.; Chen, L.; and Amer-Yahia, S. 2016. A Survey of General-Purpose Crowdsourcing Techniques. *IEEE Transactions on Knowledge and Data Engineering*, 28: 2246–2266.

Li, G.; Wang, J.; Zheng, Y.; and Franklin, M. J. 2016. Crowdsourced Data Management: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 28: 2296–2319.

Peer, E.; Brandimarte, L.; Samat, S.; and Acquisti, A. 2017. Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70: 153–163.

Aizawa, A.; Bergeron, F.; Chen, J.; Cheng, F.; Hayashi, K.; Inui, K.; Ito, H.; Kawahara, D.; Kitsuregawa, M.; Kiyomaru, H.; et al. 2020. A System for Worldwide COVID-19 Information Aggregation. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*. Online: Association for Computational Linguistics.

Tran-Thanh, L.; Chapman, A.; Rogers, A.; and Jennings, N. 2012. Knapsack based Optimal Policies for Budget-Limited Multi-Armed Bandits. *Proceedings of the National Conference on Artificial Intelligence*, 26: 1134–1140.