An Online Approach to Task Assignment and Sequencing in Expert Crowdsourcing

Ioanna Lykourentzou
Luxembourg Institute of Science and Technology
ioanna.lykourentzou@list.lu

Heinz Schmitz
Trier University of Applied Sciences, Germany
h.schmitz@hochschule-trier.de

Abstract

We introduce the problem of crowdsourcing Task Assignment and Sequencing (TAS), which adds the timeline perspective to expert crowdsourcing optimization. Subsequently we propose TAS-ONLINE, an online algorithm that aims to complete as many tasks as possible within budget and required quality without future input information regarding job release dates or worker availabilities. Results, comparing TAS-ONLINE to four typical benchmarks, show that it can achieve more completed jobs, lower flow times and higher job quality. This work has practical implications for improving the outcome of expert crowdsourcing with cost, quality and time constraints.

Introduction. As the appeal of expert crowdsourcing increases, there is a growing need to *optimize* its outputs in terms of cost, quality and timeliness. Recent studies that examine this problem (Basu Roy et al. 2015; Goel, Nikzad, and Singla 2014) typically seek to find worker allocations per task such that the worker contributions add up to a required quality threshold within a given budget. Roughly speaking, the crowdsourced tasks play the roles of multiple knapsacks with some additional concepts like domain-specific expertise and wages per worker, different models of acceptance probabilities or types of quality aggregation. However the aspect of time and the variability of workers or tasks (i.e. the dynamic aspect of the optimization problem) are not sufficiently taken into account. Optimizing for time raises the need not only for worker-to-task allocation but also for scheduling. Optimizing dynamically raises the need not only for static, index-based algorithms, but for online ones.

In this paper we introduce the problem of crowdsourcing Task Assignment and Sequencing (TAS), which adds the *timeline perspective* to the crowdsourcing optimization model: How can we find task assignments that can be rolled out in a realistic timeline, featuring unknown task release dates and worker availabilities? Overall, our three main contributions are: i) explicitly adding the *timeline perspective* to task assignment modeling in expert crowdsourcing ii) proposing an *online algorithm*, TAS-ONLINE, which tries to complete as many jobs as possible within budget and required quality, without future input information regarding job release dates or worker availabilities, and iii) illustrating through simulated and real-world experiments that TAS-ONLINE can achieve more completed jobs, lower flow times and higher quality compared to four typical benchmarks.

Task Assignment and Sequencing (TAS). The input data consist of a scheduling period of t timeslots or days, a finite set K of knowledge domains, a finite set U of crowd workers and a finite set J of jobs. Each worker has an expertise vector specifying for each domain the added quality she can bring to a job, and an availability vector of dimension t with 0-1 entries for each day. Each job belongs to a domain, requires a quality threshold, has a budget and comes with a release date. The solution comprises, for each job, a vector of dimension t with entries from $U \cup \{none\}$ indicating that a certain (or no) worker is assigned to this job on a particular day. To be feasible, a solution must respect a number of constraints: a) no worker is assigned to more than one job at a time, b) no job is assigned to more than one worker at a time, c) no worker is assigned more than once to the same job, d) no worker is scheduled on a day that she is not available, e) no job is worked on before its release date, and f) no job exceeds its budget. To assess the quality of a feasible solution we count the number of jobs that reach their quality threshold within budget and can be scheduled in a feasible way w.r.t. the above constraints. We call such jobs *completed*. TAS is a strongly NP-hard optimization problem, which combines aspects of the multiple knapsack decision problem with a unit-time openshop problem with limited machine availabilities and job release dates.

An Online Algorithm for TAS. Due to the dynamic nature of crowdsourcing it seems not realistic to consider TAS as an *offline* problem where the algorithm can be provided with the complete input at once. In fact, worker availabilities are hardly predictable and it is usually not known in advance which jobs will enter the system at what time. So the problem of task assignment and sequencing is inherently *online* and scheduling decisions have to be taken without complete input information Instead in each step the algorithm must make decisions having access only to time-dependent information from the past. The starting point of our algorithm design is the observation that in each feasible solution the worker-task assignments per timeslot form a matching between the released but incomplete jobs and the available workers. The remaining constraints can be respected by omitting edges in the underlying bipartite graph while look-

Table 1: Objective function and further performance metrics

Algorithm	absolute	% upper bound	flow time	budget usage	quality reached
RAND.	2	$0,39$	$4,77$	$88,16$	$60,62$
RAND. EGOISTIC	98	$19,03$	$3,46$	$92,05$	$90,27$
RAND. EGOISTIC FILTER	114	$22,14$	$7,37$	$52,6$	$55,77$
TAS-ONLINE	**355**	$\mathbf{68,93}$	$\mathbf{3,31}$	$\mathbf{94,35}$	$\mathbf{97,76}$
TAS-OFFLINE	411	$79,81$	$8,11$	$67,73$	$70,21$

ing at information from the past. Further we want a sequence of matchings that yields a large number of completed jobs. We thus propose a greedy approach, which computes in each step a matching such that the sum of *profits* obtained from the respective assignments for that timeslot is maximized. The resulting algorithm, TAS-ONLINE, has polynomial running time $O(t|J|^3|U|^3)$, and constitutes an online schema that allows multiple extensions, discussed in the last section.

Evaluation. We evaluate the performance of TAS-ONLINE using 4 benchmarks: 1) RANDOM builds a feasible solution randomly, without any individual worker preferences that usually appear in a self-organized system, 2) RANDOM EGOISTIC models a typical crowdsourcing environment where workers are self-appointed to tasks, trying to maximize their profit by prioritizing highest paying jobs (Rogstadius et al. 2011), 3) RANDOM EGOISTIC FILTER extends the previous by restricting the offered jobs depending on worker expertise, modeling current crowd worker screening practices (Downs et al. 2010), and finally for reasons of comparison 4) TAS-OFFLINE, an offline version of TAS-ONLINE, which has access to the complete input at once.

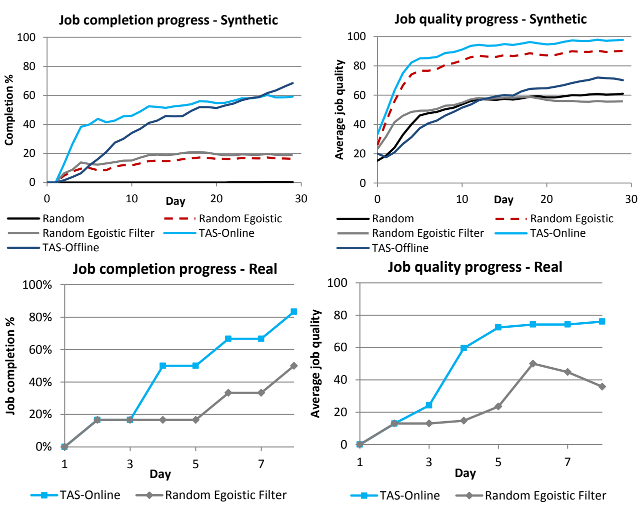

Figure 1: Synthetic (a,b) and Real-world (c,d) experiments

Synthetic Instance Experiment. We first experiment with synthetic data, generated using the experimental distributions reported in (Basu Roy et al. 2015), where AMT workers worked on the complex task of news writing. We simulate a timeline of 30 days, during which 1000 workers (re)entered the system and 600 jobs were requested, belonging to 10 knowledge domains. In terms of the objective function (Table 1) we observe that TAS-ONLINE does not reach the number of completed jobs of the offline algorithm, but it is significantly better than the other online benchmarks. We further compare the algorithms w.r.t. flow time length (time between release date r_j and latest assigned worker for j), average job quality and budget utilization. We observe that TAS-ONLINE achieves faster job accomplishment (shorter flow times) and higher quality than all its competitors, although it consumes most of its budget due to its greedy nature. We also observe (Figure 1 (a,b)) that TAS-ONLINE achieves consistently higher job completion rates, beaten only towards the end of the scheduling period by TAS-OFFLINE possibly due to the lookahead mechanism of the latter that takes the timeline end into account. Last, TAS-ONLINE consistently achieves higher average job quality throughout the scheduling period.

Real-World Experiment. Next, we test the algorithm in real-world conditions. The task used was collaborative news article writing, with workers building on each other's content sequentially, on two topics: "FIFA 2015 corruption scandal" and "Self-driving cars". The scheduling period was 8 days and the worker pool was 60 CrowdFlower workers. After recording worker expertise (using two multiple choice tests) and required wage, we split workers randomly into 2 pools (one for the benchmark RANDOM EGOISTIC FILTER and one for TAS-ONLINE). We created 6 Google documents (3 per topic) corresponding to the jobs, and set up their cost/quality thresholds and release dates similarly to the synthetic experiments. Each day a worker was assigned to a job and then the job's current quality was evaluated by 50 independent workers, until the job passed its quality threshold or exhausted its budget. At the end of the scheduling period the benchmark accomplished 3/6 jobs and TAS-ONLINE 5/6. TAS-ONLINE also achieved higher average job quality and completion percentages (Figure 1 (c,d)). These results are all in line with the synthetic data results.

Conclusion and Future Extensions. In this paper we present TAS, a model that adds the timeline and online perspective to task assignment optimization in expert crowdsourcing. Using a greedy scheduling algorithm, TAS-ONLINE, we show that optimization under this model can significantly improve expert crowdsourcing performance. Future extensions include adding budget flexibility, job and/or worker prioritizing, worker performance variability over time, and multiple assignments per worker/timeslot.

References

Basu Roy, S.; Lykourentzou, I.; Thirumuruganathan, S.; Amer-Yahia, S.; and Das, G. 2015. Task assignment optimization in knowledge-intensive crowdsourcing. *VLDB Journal* 1–25.

Downs, J. S.; Holbrook, M. B.; Sheng, S.; and Cranor, L. F. 2010. Are your participants gaming the system?: Screening mechanical turk workers. CHI '10, Atlanta, April 10-15.

Goel, G.; Nikzad, A.; and Singla, A. 2014. Allocating tasks to workers with matching constraints: Truthful mechanisms for crowdsourcing markets. WWW'14, Seoul, April 7-11.

Rogstadius, J.; Kostakos, V.; Kittur, A.; Smus, B.; Laredo, J.; and Vukovic, M. 2011. An assessment of intrinsic and extrinsic motivation on task performance in crowdsourcing markets. ICWSM'11, Barcelona, July 17-21.