

Statistical Decision Making for Optimal Budget Allocation in Crowd Labeling

Xi Chen

*Stern School of Business
New York University
New York, New York, 10012, USA*

XICHEN@NYU.EDU

Qihang Lin

*Tippie College of Business
University of Iowa
Iowa City, Iowa, 52242, USA*

QIHANG-LIN@UIOWA.EDU

Dengyong Zhou

*Microsoft Research
Redmond, Washington, 98052, USA*

DENGYONG.ZHOU@MICROSOFT.COM

Editor: Yuan (Alan) Qi

Abstract

It has become increasingly popular to obtain machine learning labels through commercial crowdsourcing services. The crowdsourcing workers or annotators are paid for each label they provide, but the task requester usually has only a limited amount of the budget. Since the data instances have different levels of labeling difficulty and the workers have different reliability for the labeling task, it is desirable to wisely allocate the budget among all the instances and workers such that the overall labeling quality is maximized. In this paper, we formulate the budget allocation problem as a Bayesian Markov decision process (MDP), which simultaneously conducts learning and decision making. The optimal allocation policy can be obtained by using the dynamic programming (DP) recurrence. However, DP quickly becomes computationally intractable when the size of the problem increases. To solve this challenge, we propose a computationally efficient approximate policy which is called optimistic knowledge gradient. Our method applies to both pull crowdsourcing marketplaces with homogeneous workers and push marketplaces with heterogeneous workers. It can also incorporate the contextual information of instances when they are available. The experiments on both simulated and real data show that our policy achieves a higher labeling quality than other existing policies at the same budget level.

Keywords: crowdsourcing, budget allocation, Markov decision process, dynamic programming, optimistic knowledge gradient

1. Introduction

In many real applications, data are usually collected without any innate label. For example, a digital camera will not automatically tag a picture as a portrait or a landscape. A traditional approach for data labeling is to hire a small group of experts to provide labels for the entire set of data. However, for large-scale data, such an approach becomes inefficient and very costly. Thanks to the advent of many online crowdsourcing services, e.g., Amazon

Mechanical Turk, a much more efficient way is to post unlabeled data to a crowdsourcing marketplace, where a big crowd of low-paid workers can be hired instantaneously to perform labeling tasks.

Despite its high efficiency and immediate availability, crowd labeling raises many new challenges. Since labeling tasks are tedious and workers are usually non-experts, labels generated by the crowd suffer from low quality. As a remedy, most crowdsourcing services resort to labeling redundancy to reduce the labeling noise, which is achieved by collecting multiple labels from different workers for each data instance. In particular, a crowd labeling process can be described as a two phase procedure:

1. In the first phase, unlabeled data instances are assigned to a crowd of workers and multiple raw labels are collected for each data instance.
2. In the second phase, for each data instance, one aggregates the collected raw labels to infer its true label.

In principle, more raw labels will lead to a higher chance of recovering the true label. However, each raw label comes with a cost: the requester has to pay workers pre-specified monetary reward for each label they provide, usually, regardless of the label’s correctness. For example, a worker typically earns 10 cents by categorizing a website as porn or not. In practice, the requester has only a limited amount of budget which essentially restricts the total number of raw labels that he/she can collect. This raises a challenging question central in crowd labeling: *What is the best way to allocate the budget among data instances and workers so that the overall accuracy of aggregated labels is maximized ?*

The most important factors that decide how to allocate the budget are the intrinsic characteristics of data instances and workers: *labeling difficulty/ambiguity for each data instance* and *reliability/quality of each worker*. In particular, an instance is less ambiguous if its label can be decided based on the common knowledge and a vast majority of reliable workers will provide the same label for it. In principle, we should avoid spending too much budget on those easy instances since excessive raw labels will not bring much additional information. In contrast, for an ambiguous instance which falls near the boundary of categories, even those reliable workers will still disagree with each other and generate inconsistent labels. For those ambiguous instances, we are facing a challenging decision problem on how much budget that we should spend on them. On one hand, it is worth to collect more labels to boost the accuracy of the aggregate label. On the other hand, since our goal is to maximize the *overall* labeling accuracy, when the budget is limited, we should simply put those few highly ambiguous instances aside to save budget for labeling less difficult instances. In addition to the ambiguity of data instances, the other important factor is the reliability of each worker and, undoubtedly, it is desirable to assign more instances to those reliable workers. Despite their importance in deciding how to allocate the budget, both the data ambiguity and workers’ reliability are unknown parameters at the beginning and need to be updated based on the stream of collected raw labels in an online fashion. This further suggests that the budget allocation policy should be dynamic and simultaneously conduct parameter estimation and decision making.

To search for an optimal budget allocation policy, we model the data ambiguity and workers’ reliability using two sets of random variables drawn from known prior distributions. Then, we formulate the problem into a finite-horizon Bayesian Markov Decision

Process (MDP) (Puterman, 2005), whose state variables are the posterior distributions of these variables, which are updated by each new label. Here, the Bayesian setting is necessary. We will show that an optimal policy only exists in the Bayesian setting. Using the MDP formulation, the optimal budget allocation policy for any finite budget level can be readily obtained via the dynamic programming (DP). However, DP is computationally intractable for large-scale problems since the size of the state space grows exponentially in budget level. The existing widely-used approximate policies, such as approximate Gittins index rule (Gittins, 1989) or knowledge gradient (KG) (Gupta and Miescke, 1996; Frazier et al., 2008), either has a high computational cost or poor performance in our problem. In this paper, we propose a new policy, called *optimistic knowledge gradient (Opt-KG)*. In particular, the Opt-KG policy dynamically chooses the next instance-worker pair based on the optimistic outcome of the marginal improvement on the accuracy, which is a function of state variables. We further propose a more general Opt-KG policy using the conditional value-at-risk measure (Rockafellar and Uryasev, 2002). The Opt-KG is computationally efficient, achieves superior empirical performance and has some asymptotic theoretical guarantees.

To better present the main idea of our MDP formulation and the Opt-KG policy, we start from the binary labeling task (i.e., providing the category, either positive or negative, for each instance). We first consider the *pull marketplace* (e.g., Amazon Mechanical Turk or Galaxy Zoo), where the labeling requester can only post instances to the general worker pool with either anonymous or transient workers, but cannot assign to an identified worker. In a pull marketplace, workers are typically treated as *homogeneous* and one models the entire worker pool instead of each individual worker. We further assume that workers are fully reliable (or noiseless) such that the chance that they make an error only depend on instances' own ambiguity. At a first glance, such an assumption may seem oversimplified. In fact, it turns out that the budget-optimal crowd labeling under such an assumption has been highly non-trivial. We formulate this problem into a Bayesian MDP and propose the computational efficient Opt-KG policy. We further prove that the Opt-KG policy in such a setting is asymptotically consistent, that is, when the budget goes to infinity, the accuracy converges to 100% almost surely.

Then, we extend the MDP formulation to deal with *push marketplaces* with *heterogeneous* workers. In a push marketplace (e.g., data annotation team in Microsoft Bing group), once an instance is allocated to an identified worker, the worker is required to finish the instance in a short period of time. Based on the previous model for fully reliable workers, we further introduce another set of parameters to characterize workers' reliability. Then our decision process simultaneously selects the next instance to label and the next worker for labeling the instance according to the optimistic knowledge gradient policy. In fact, the proposed MDP framework is so flexible that we can further extend it to incorporate contextual information of instances whenever they are available (e.g., as in web search and advertising applications discussed in Li et al., 2010) and to handle multi-class labeling.

In summary, the main contribution of the paper consists of the three folds: (1) we formulate the budget allocation in crowd labeling into a MDP and characterize the *optimal* policy using DP; (2) computationally, we propose an efficient approximate policy, optimistic knowledge gradient; (3) the proposed MDP framework can be used as a general framework to address various budget allocation problems in crowdsourcing (e.g., rating and ranking tasks).

The rest of this paper is organized as follows. In Section 2, we first present the modeling of budget allocation process for binary labeling tasks with fully reliable workers and motivate our Bayesian modeling. In Section 3, we present the Bayesian MDP and the optimal policy via DP. In Section 4, we propose a computationally efficient approximate policy, Opt-KG. In Section 5, we extend our MDP to model heterogeneous workers with different reliability. In Section 6, we present other important extensions, including incorporating contextual information and multi-class labeling. In Section 7, we discuss the related works. In Section 8, we present numerical results on both simulated and real data sets, followed by conclusions in Section 9.

2. Binary Labeling with Homogeneous Noiseless Workers

We first consider the budget allocation problem in a pull marketplace with homogeneous noiseless workers for binary labeling tasks. We note that such a simplification is important for investigating this problem, since the incorporation of workers' reliability and extensions to multiple categories become rather straightforward once this problem is correctly modeled (see Section 5 and 6).

Suppose that there are K instances and each one is associated with a latent true label $Z_i \in \{-1, 1\}$ for $1 \leq i \leq K$. Our goal is to infer the set of positive instances, denoted by $H^* = \{i : Z_i = 1\}$. Here, we assume that the homogeneous worker pool is *fully reliable* or *noiseless*. We note that it does not mean that each worker knows the true label Z_i . Instead, it means that fully reliable workers will do their best to make judgments but their labels may be still incorrect due to the instance's ambiguity. Further, we model the *labeling difficulty/ambiguity* of each instance by a latent soft-label θ_i , which can be interpreted as the percentage of workers in the homogeneous noiseless crowd who will label the i -th instance as positive. In other words, if we randomly choose a worker from a large crowd of fully reliable workers, we will receive a positive label for the i -th instance with probability θ_i and a negative label with probability $1 - \theta_i$. In general, we assume the crowd is large enough so that the value of θ_i can be any value in $[0, 1]$. To see how θ_i characterizes the labeling difficulty of the i -th instance, we consider a concrete example where a worker is asked to label a person as adult (positive) or not (negative) based on the photo of that person. If the person is more than 25 years old, most likely, the corresponding θ_i will be close to 1, generating positive labels consistently. On the other hand, if the person is younger than 15, she may be labeled as negative by almost all the reliable workers since θ_i is close to 0. In both of this cases, we regard the instance (person) easy to label since Z_i can be inferred with a high accuracy based on only a few raw labels. On the contrary, for a person is one or two years below or above 18, the θ_i is near 0.5 and the numbers of positive and negative labels become relatively comparable so that the corresponding labeling task is very difficult. Given the definition of soft labels, we further make the following assumption:

Assumption 1 *We assume that the soft-label θ_i is consistent with the true label in the sense that $Z_i = 1$ if and only if $\theta_i \geq 0.5$, i.e., the majority of the crowd are correct, and hence $H^* = \{i : \theta_i \geq 0.5\}$.*

Given the total budget, denoted by T , we suppose that each label costs one unit of budget. As discussed in the introduction, the crowd labeling has two phases. The first

Instance	1st round label	2nd round label
Instance 1 (θ_1)	1	1
Instance 2 (θ_2)	1	-1
Instance 3 (θ_3)	1	

Table 1: Toy example with 3 instances and 5 collected labels. Instance 1 has two positive labels, instance 2 has one positive and one negative label, and instance 3 has only one positive label. The question is that, given only one more labeling chance, which instance should be chosen to label?

phase is the *budget allocation phase*, which is a dynamic decision process with T stages. In each stage $0 \leq t \leq T - 1$, an instance $i_t \in \mathcal{A} = \{1, \dots, K\}$ is selected based on the historical labeling results. Once i_t is selected, it will be labeled by a random worker from the homogeneous noiseless worker pool. According to the definition of θ_{i_t} , the label received, denoted by $y_{i_t} \in \{-1, 1\}$, will follow the Bernoulli distribution with the parameter θ_{i_t} :

$$\Pr(y_{i_t} = 1) = \theta_{i_t} \quad \text{and} \quad \Pr(y_{i_t} = -1) = 1 - \theta_{i_t}. \quad (1)$$

We note that, at this moment, all workers are assumed to be homogeneous and noiseless so that y_{i_t} only depends on θ_{i_t} but not on which worker provides the label. Therefore, it is suffice for the decision maker (e.g., requester or crowdsourcing service) to select the instance in each stage instead of an instance-worker pair.

The second phase is the *label aggregation phase*. When the budget is exhausted, the decision maker needs to infer true labels $\{Z_i\}_{i=1}^n$ by aggregating all the collected labels. According to Assumption 1, it is equivalent to infer the set of positive instances whose $\theta_i \geq 0.5$. Let H_T be the estimated positive set. The final overall accuracy is measured by $|H_T \cap H^*| + |(H_T)^c \cap (H^*)^c|$, the size of the mutual overlap between H^* and H_T .

Our goal is to determine the *optimal allocation policy*, (i_0, \dots, i_{T-1}) , so that overall accuracy is maximized. Here, a natural question to ask is whether the *optimal* allocation policy exists and what assumptions do we need for the existence of the optimal policy. To answer this question, we provide a concrete example, which motivates our Bayesian modeling.

2.1 Why We Need a Bayesian Modeling

Let us check a toy example with 3 instances and 5 collected labels (see Table 1). We assume that the workers are homogeneous noiseless and the label aggregation is performed by the majority vote rule. *Now if we only have the budget to get one more label, which instance should be chosen to label?* It is obvious that we should not put the remaining budget on the first instance since we are relatively more confident on what its true label should be. Thus, the problem becomes how to choose between the second and third instances. In what follows, we shall show that *there is no optimal policy under the frequentist setting*. To be more explicit, the optimal policy leads to the expected accuracy which is at least as good as that of all other policies for any values of $\{\theta_i\}_{i=1}^n$.

	Current Accuracy	$y = 1$	$y = -1$	Expected Accuracy	Improvement
$\theta_1 > 0.5$	1	1	1	1	0
$\theta_1 < 0.5$	0	0	0	0	0
$\theta_2 > 0.5$	0.5	1	0	θ_2	$\theta_2 - 0.5 > 0$
$\theta_2 < 0.5$	0.5	0	1	$1 - \theta_2$	$0.5 - \theta_2 > 0$
$\theta_3 > 0.5$	1	1	0.5	$\theta_3 + 0.5(1 - \theta_3)$	$0.5(\theta_3 - 1) < 0$
$\theta_3 < 0.5$	0	0	0.5	$0.5(1 - \theta_3)$	$0.5(1 - \theta_3) > 0$

Table 2: Expected improvements in accuracy for collecting an extra label, i.e., the expected accuracy of obtaining one more label minus the current expected accuracy. The 3rd and 4th columns contain the accuracies with the next label being 1 and -1 . The 5th is the expected accuracy which is computed by taking θ times the 3rd column plus $(1 - \theta)$ times the 4th. The last column contains the expected improvements which is computed by taking the difference between the 5th and 2nd columns.

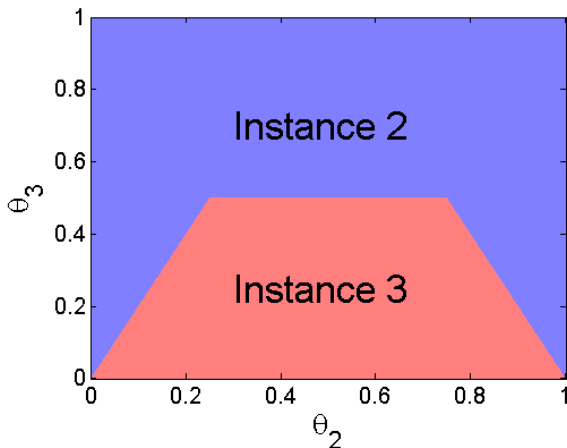


Figure 1: Decision Boundary.

Let us compute the expected improvement in accuracy in terms of the frequentist risk in Table 2. We assume that $\theta_i \neq 0.5$ and if the number of 1 and -1 labels are the same for an instance, the accuracy is 0.5 based on a random guess. From Table 2, we should not label the first instance since the improvement is always 0. This coincides with our intuition. When $\max(\theta_2 - 0.5, 0.5 - \theta_2) > 0.5(1 - \theta_3)$ or $\theta_3 > 0.5$, which corresponds to the blue region in Figure 1, we should choose to label the second instance. Otherwise, we should ask the label for the third one. Since the true value of θ_2 and θ_3 are unknown, an optimal policy does not exist under the frequentist paradigm. Further, it will be difficult to estimate θ_2 and θ_3 accurately when the budget is very limited.

In contrast, in a Bayesian setting with prior distribution on each θ_i , the optimal policy is defined as the policy which leads to the highest expected accuracy under the given prior

instead of for any possible values of $\{\theta_i\}_{i=1}^n$. Therefore, we can optimally determine the next instance to label by taking another expectation over the distribution of θ_i . In this paper, we adopt the Bayesian modeling to formulate the budget allocation problem in crowd labeling.

3. Bayesian MDP and Optimal Policy

In this section, we first introduce the Bayesian MDP for modeling the dynamic budget allocation process and then provide the optimal allocation policy using dynamic programming.

3.1 Bayesian Modeling

We assume that each θ_i is drawn from a known Beta prior $\text{Beta}(a_i^0, b_i^0)$. Beta is a rich family of distributions in the sense that it exhibits a fairly wide variety of shapes on the domain of θ_i , i.e., the unit interval $[0, 1]$. For presentation simplicity, instead of considering a full Bayesian model with hyper-priors on a_i^0 and b_i^0 , we fix a_i^0 and b_i^0 at the beginning. In practice, if the budget is sufficient, one can first label each instance equally many times to pre-estimate $\{a_i^0, b_i^0\}_{i=1}^K$ before the dynamic labeling procedure is invoked. Otherwise, when there is no prior knowledge, we can simply assume $a_i^0 = b_i^0 = 1$ so that the prior is a uniform distribution. According to our simulated experimental results in Section 8.1.2, uniform prior works reasonably well unless the data is highly skewed in terms of class distribution. Other commonly used uninformative priors such as Jeffreys prior or reference prior ($\text{Beta}(1/2, 1/2)$) or Haldane prior ($\text{Beta}(0, 0)$) can also be adopted (see Robert, 2007 for more on uninformative priors). Choices of prior distributions are discussed in more details in Section 4.2.

At each stage t with $\text{Beta}(a_i^t, b_i^t)$ as the current posterior distribution for θ_i , we make a decision by choosing an instance $i_t \in \mathcal{A} = \{1, \dots, K\}$ and acquire its label $y_{i_t} \sim \text{Bernoulli}(\theta_{i_t})$. Here \mathcal{A} denotes the *action set*. By the fact that Beta is the conjugate prior of the Bernoulli, the posterior of θ_{i_t} in the stage $t + 1$ will be updated as:

$$\text{Beta}(a_{i_t}^{t+1}, b_{i_t}^{t+1}) = \begin{cases} \text{Beta}(a_{i_t}^t + 1, b_{i_t}^t) & \text{if } y_{i_t} = 1; \\ \text{Beta}(a_{i_t}^t, b_{i_t}^t + 1) & \text{if } y_{i_t} = -1. \end{cases}$$

We put $\{a_i^t, b_i^t\}_{i=1}^K$ into a $K \times 2$ matrix S^t , called a *state matrix*, and let $S_i^t = (a_i^t, b_i^t)$ be the i -th row of S^t . The update of the state matrix can be written in a more compact form:

$$S^{t+1} = \begin{cases} S^t + (\mathbf{e}_{i_t}, \mathbf{0}) & \text{if } y_{i_t} = 1; \\ S^t + (\mathbf{0}, \mathbf{e}_{i_t}) & \text{if } y_{i_t} = -1, \end{cases} \quad (2)$$

where \mathbf{e}_{i_t} is a $K \times 1$ vector with 1 at the i_t -th entry and 0 at all other entries. As we can see, $\{S^t\}$ is a Markovian process because S^{t+1} is completely determined by the current state S^t , the action i_t and the obtained label y_{i_t} . It is easy to calculate the *state transition probability* $\Pr(y_{i_t} | S^t, i_t)$, which is the posterior probability that we are in the next state S^{t+1} if we choose i_t to be label in the current state S^t :

$$\Pr(y_{i_t} = 1 | S^t, i_t) = \mathbb{E}(\theta_{i_t} | S^t) = \frac{a_{i_t}^t}{a_{i_t}^t + b_{i_t}^t} \quad \text{and} \quad \Pr(y_{i_t} = -1 | S^t, i_t) = \frac{b_{i_t}^t}{a_{i_t}^t + b_{i_t}^t}. \quad (3)$$

Given this labeling process, the budget allocation policy is defined as a sequence of decisions: $\pi = (i_0, \dots, i_{T-1})$. Here, we require decisions depend only upon the previous information. To make this more formal, we define a filtration $\{\mathcal{F}_t\}_{t=0}^T$, where \mathcal{F}_t is the information collected until the stage $t - 1$. More precisely, \mathcal{F}_t is the σ -algebra generated by the sample path $(i_0, y_{i_0}, \dots, i_{t-1}, y_{i_{t-1}})$. We require the action i_t is determined based on the historical labeling results up to the stage $t - 1$, i.e., i_t is \mathcal{F}_t -measurable.

3.2 Inference About True Labels

As described in Section 2, the budget allocation process has two phases: the dynamic budget allocation phase and the label aggregation phase. Since the goal of the dynamic budget allocation in the first phase is to maximize the accuracy of aggregated labels in the second phase, we first present how to infer the true label via label aggregation in the second phase.

When the decision process terminates at the stage T , we need to determine a positive set H_T to maximize the *conditional* expected accuracy conditioning on \mathcal{F}_T , which corresponds to minimizing the posterior risk:

$$H_T = \arg \max_{H \subset \{1, \dots, K\}} \mathbb{E} \left(\sum_{i=1}^K (\mathbf{1}(i \in H) \cdot \mathbf{1}(i \in H^*) + \mathbf{1}(i \notin H) \cdot \mathbf{1}(i \notin H^*)) \middle| \mathcal{F}_T \right), \quad (4)$$

where $\mathbf{1}(A)$ is the indicator function, which takes the value 1 if the event A is true and 0 otherwise. The term inside expectation in (4) is the binary labeling accuracy which can also be written as $|H \cap H^*| + |H^c \cap (H^*)^c|$.

We first observe that, for $0 \leq t \leq T$, the conditional distribution $\theta_i | \mathcal{F}_t$ is exactly the posterior distribution $\text{Beta}(a_i^t, b_i^t)$, which depends on the historical sampling results only through $S_i^t = (a_i^t, b_i^t)$. Hence, we define

$$I(a, b) \doteq \Pr(\theta \geq 0.5 | \theta \sim \text{Beta}(a, b)), \quad (5)$$

$$P_i^t \doteq \Pr(i \in H^* | \mathcal{F}_t) = \Pr(\theta_i \geq 0.5 | \mathcal{F}_t) = \Pr(\theta_i \geq 0.5 | S_i^t) = I(a_i^t, b_i^t). \quad (6)$$

As shown in Xie and Frazier (2013), the optimal positive set H_T can be determined by the Bayes decision rule as follows.

Proposition 2 $H_T = \{i : \Pr(i \in H^* | \mathcal{F}_T) \geq 0.5\} = \{i : P_i^T \geq 0.5\}$ solves (4).

The proof of Proposition 2 is given in the appendix for completeness.

With Proposition 2 in place, we plug the optimal positive set H_T into the right hand side of (4) and the conditional expected accuracy given \mathcal{F}_T can be simplified as:

$$\mathbb{E} \left(\sum_{i=1}^K (\mathbf{1}(i \in H_T) \cdot \mathbf{1}(i \in H^*) + \mathbf{1}(i \notin H_T) \cdot \mathbf{1}(i \notin H^*)) \middle| \mathcal{F}_T \right) = \sum_{i=1}^K h(P_i^T), \quad (7)$$

where $h(x) \doteq \max(x, 1 - x)$. We also note that P_i^T provides not only the estimated label for the i -th instance but also how confident the estimated label is correct. According to the next corollary with the proof in the appendix, we show that the optimal H_T is constructed based on a refined *majority vote* rule which incorporates the prior information.

Corollary 3 $I(a, b) > 0.5$ if and only if $a > b$ and $I(a, b) = 0.5$ if and only if $a = b$. Therefore, $H_T = \{i : a_i^T \geq b_i^T\}$ solves (4).

By viewing a_i^0 and b_i^0 as pseudo-counts of 1s and -1 s at the initial stage, the parameters a_i^T and b_i^T are the total counts of 1s and -1 s. The estimated positive set $H_T = \{i : a_i^T \geq b_i^T\}$ consists of instances with more (or equal) counts of 1s than that of -1 s. When $a_i^0 = b_i^0$, H_T is constructed exactly according to the vanilla *majority vote* rule.

To find the optimal allocation policy which maximizes the expected accuracy, we need to solve the following optimization problem:

$$\begin{aligned} V(S^0) &\doteq \sup_{\pi} \mathbb{E}^{\pi} \left[\mathbb{E} \left(\sum_{i=1}^K (\mathbf{1}(i \in H_T) \cdot \mathbf{1}(i \in H^*) + \mathbf{1}(i \notin H_T) \cdot \mathbf{1}(i \notin H^*)) \middle| \mathcal{F}_T \right) \right] \\ &= \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{i=1}^K h(P_i^T) \right), \end{aligned} \quad (8)$$

where \mathbb{E}^{π} represents the expectation taken over the sample paths $(i_0, y_{i_0}, \dots, i_{T-1}, y_{i_{T-1}})$ generated by a policy π . The second equality is due to Proposition 2 and $V(S^0)$ is called value function at the initial state S^0 . The optimal policy π^* is any policy π that attains the supremum in (8).

3.3 Markov Decision Process

The optimization problem in (8) is essentially a Bayesian multi-armed bandit (MAB) problem, where each instance corresponds to an arm and the decision is which instance/arm to be sampled next. However, it is different from the classical MAB problem (Auer et al., 2002; Bubeck and Cesa-Bianchi, 2012), which assumes that each sample of an arm yields independent and identically distributed (*i.i.d.*) reward according to some unknown distribution associated with that arm. Given the total budget T , the goal is to determine a sequential allocation policy so that the collected rewards can be maximized. We contrast this problem with our problem: instead of collecting intermediate independent rewards on the fly, our objective in (8) merely involves the final “reward”, i.e., overall labeling accuracy, which is only available at the final stage when the budget runs out. Although there is no intermediate reward in our problem, we can still decompose the final expected accuracy into sum of *stage-wise rewards* using the technique from Xie and Frazier (2013), which further leads to our MDP formulation. Since these *stage-wise rewards* are artificially created, they are no longer *i.i.d.* for each instance. We also note that the problem in Xie and Frazier (2013) is an infinite-horizon one which optimizes the stopping time while our problem is *finite-horizon* since the decision process must be stopped at the stage T .

Proposition 4 Define the stage-wise expected reward as:

$$R(S^t, i_t) = \mathbb{E} \left(\sum_{i=1}^K h(P_i^{t+1}) - \sum_{i=1}^K h(P_i^t) \middle| S^t, i_t \right) = \mathbb{E} (h(P_{i_t}^{t+1}) - h(P_{i_t}^t) | S^t, i_t), \quad (9)$$

then the value function (8) becomes:

$$V(S^0) = G_0(S^0) + \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{t=0}^{T-1} R(S^t, i_t) \right), \quad (10)$$

where $G_0(S^0) = \sum_{i=1}^K h(P_i^0)$ and the optimal policy π^* is any policy π that attains the supremum.

The proof of Proposition 4 is presented in the appendix. In fact, the stage-wise reward in (9) has a straightforward interpretation. According to (8), the term $\sum_{i=1}^K h(P_i^t)$ is the expected accuracy at the t -th stage. The stage-wise reward $R(S^t, i_t)$ takes the form of the difference between the expected accuracy at the $(t+1)$ -stage and the t -th stage, i.e., the expected gain in accuracy for collecting another label for the i_t -th instance. The second equality in (9) holds simply because: only the i_t -th instance receives the new label and the corresponding $P_{i_t}^t$ changes while all other P_i^t remain the same. Since the expected reward (9) only depends on $S_{i_t}^t = (a_{i_t}^t, b_{i_t}^t)$, we write

$$R(S^t, i_t) = R(S_{i_t}^t) = R(a_{i_t}^t, b_{i_t}^t), \quad (11)$$

and use them interchangeably. The function $R(a, b)$ with two parameters a and b has an analytical representation as follows. For any state (a, b) of a single instance, the reward of getting a label 1 and a label -1 are:

$$R_1(a, b) = h(I(a+1, b)) - h(I(a, b)), \quad (12)$$

$$R_2(a, b) = h(I(a, b+1)) - h(I(a, b)). \quad (13)$$

The expected reward takes the following form:

$$R(a, b) = p_1 R_1 + p_2 R_2, \quad (14)$$

where $p_1 = \frac{a}{a+b}$ and $p_2 = \frac{b}{a+b}$ are the transition probabilities in (3).

With Proposition 4, the maximization problem (8) is formulated as a T -stage *Markov Decision Process* (MDP) as in (10), which is associated with a tuple:

$$\{T, \{\mathcal{S}^t\}, \mathcal{A}, \Pr(y_{i_t}|S^t, i_t), R(S^t, i_t)\}.$$

Here, the state space at the stage t , \mathcal{S}^t , is all possible states that can be reached at t . Once we collect a label y_{i_t} , one element in S^t (either $a_{i_t}^t$ or $b_{i_t}^t$) will add one. Therefore, we have

$$\mathcal{S}^t = \left\{ \{a_i^t, b_i^t\}_{i=1}^K : a_i^t \geq a_i^0, b_i^t \geq b_i^0, \sum_{i=1}^K (a_i^t - a_i^0) + (b_i^t - b_i^0) = t \right\}. \quad (15)$$

The action space is the set of instances that could be labeled next: $\mathcal{A} = \{1, \dots, K\}$. The transition probability $\Pr(y_{i_t}|S^t, i_t)$ is defined in (3) and the expected reward at each stage $R(S^t, i_t)$ is defined in (9).

Remark 5 We can also view Proposition 4 as a consequence of applying the reward shaping technique (Ng et al., 1999) to the original problem (8). In fact, we can add an artificial

Instance i	S_i^{T-1}	p_1	p_2	$R_1(S_i^{T-1})$	$R_2(S_i^{T-1})$	$R(S^{T-1}, i) = R(S_i^{T-1})$
1	(3,1)	$\frac{3}{4}$	$\frac{1}{4}$	0.0625	-0.1875	$\frac{3}{4} \cdot (0.0625) + \frac{1}{4} \cdot (-0.1875) = 0$
2	(2,2)	$\frac{1}{2}$	$\frac{1}{2}$	0.1875	0.1875	$\frac{1}{2} \cdot (0.1875) + \frac{1}{2} \cdot (0.1875) = 0.1875$
3	(2,1)	$\frac{2}{3}$	$\frac{1}{3}$	0.1250	-0.2500	$\frac{2}{3} \cdot (0.1250) + \frac{1}{3} \cdot (-0.2500) = 0$

Table 3: Calculation of the expected reward for the toy example in Table 1 according to (12), (13) and (14).

absorbing state, named S_{obs} , to the original state space (15) and assume that, when the budget allocation process finishes, the state must transit one more time to reach S_{obs} regardless of which action is taken. Hence, the original problem (8) becomes a MDP that generates a zero transition reward until the state enters S_{obs} where the transition reward is $\sum_{i=1}^K h(P_i^T)$. Then, we define a potential-based shaping function (Ng et al., 1999) over this extended state space as $\Phi(S^t) = \sum_{i=1}^K h(P_i^t)$ for $S^t \in \mathcal{S}^t$ and $\Phi(S_{obs}) = 0$. After this, (4) can be viewed as a new MDP whose transition reward equals that of (8) plus the shaping-reward function $\Phi(S') - \Phi(S)$ when the state transits from S to S' . According to Theorem 1 in Ng et al. (1999), (4) and (8) have the same optimal policy. This provides an alternative justification for Proposition 4.

3.4 Optimal Policy via DP

With the MDP in place, we can apply the dynamic programming (DP) algorithm (a.k.a. backward induction) (Puterman, 2005) to compute the optimal policy:

1. Set $V_{T-1}(S^{T-1}) = \max_{i \in \{1, \dots, K\}} R(S^{T-1}, i)$ for all possible states $S^{T-1} \in \mathcal{S}^{T-1}$. The optimal decision $i_{T-1}^*(S^{T-1})$ is the decision i that achieves the maximum when the state is S^{T-1} .
2. Iterate for $t = T - 2, \dots, 0$, compute the $V_t(S^t)$ for all possible $S^t \in \mathcal{S}^t$ using the Bellman equation:

$$\begin{aligned}
 & V_t(S^t) \\
 &= \max_i \left(R(S^t, i) + \Pr(y_i = 1 | S^t, i) V_{t+1}(S^t + (\mathbf{e}_i, \mathbf{0})) + \Pr(y_i = -1 | S^t, i) V_{t+1}(S^t + (\mathbf{0}, \mathbf{e}_i)) \right),
 \end{aligned}$$

and $i_t^*(S^t)$ is the i that achieves the maximum.

The optimal policy $\pi^* = (i_0^*, \dots, i_T^*)$. For an illustration purpose, we use DP to calculate the optimal instance to be labeled next in the toy example in Section 2.1 under the uniform prior $B(1, 1)$ for all θ_i . Since we assume that there is only one labeling chance remaining, which corresponds to the last stage of DP, we should choose the instance $i_{T-1}^*(S^{T-1}) = \arg \max_{i \in \{1, \dots, K\}} R(S^{T-1}, i)$. According to the calculation in Table 3, there is a unique optimal instance for labeling, which is the second instance.

Although DP finds the optimal policy, its computation is intractable since the size of the state space $|\mathcal{S}^t|$ grows exponentially in t according to (15). Therefore, we need to develop a computationally efficient approximate policy, which is the goal of the next section.

4. Approximate Policies

Since DP is computationally intractable, approximate policies are needed for large-scale applications. The simplest policy is the uniform sampling (a.k.a, pure exploration), i.e., we choose the next instance uniformly and independently at random: $i_t \sim \text{Uniform}(1, \dots, K)$. However, this policy does not explore any structure of the problem.

With the decomposed reward function, our problem is essentially a finite-horizon Bayesian MAB problem. Gittins (1989) showed that Gittins index policy is optimal for infinite-horizon MAB with the discounted reward. It has been applied to the infinite-horizon version of problem (10) in Xie and Frazier (2013). Since our problem is finite-horizon, Gittins index is no longer optimal while it can still provide us a good heuristic index rule. However, the computational cost of Gittins index is very high: the state-of-art-method proposed by Nino-Mora (2011) requires $O(T^6)$ time and space complexity.

A computationally more attractive policy is the knowledge gradient (KG) (Gupta and Miescke, 1996; Frazier et al., 2008). It is essentially a single-step look-ahead policy, which greedily selects the next instance with the largest expected reward:

$$i_t = \arg \max_{i \in \{1, \dots, K\}} \left(R(a_i^t, b_i^t) \doteq \frac{a_i^t}{a_i^t + b_i^t} R_1(a_i^t, b_i^t) + \frac{b_i^t}{a_i^t + b_i^t} R_2(a_i^t, b_i^t) \right). \quad (16)$$

As we can see, this policy corresponds to the last stage in DP and hence KG policy is optimal if only one labeling chance is remaining.

When there is a tie, if we select the smallest index i , the policy is referred to *deterministic KG* while if we randomly break the tie, the policy is referred to *randomized KG*. Although KG has been successfully applied to many MDP problems (Powell, 2007), it will fail in our problem as shown in the next proposition with the proof in the appendix.

Proposition 6 *Assuming that a_i^0 and b_i^0 are positive integers and letting $\mathcal{E} = \{i : a_i^0 = b_i^0\}$, then the deterministic KG policy will acquire one label for each instance in \mathcal{E} and then consistently obtain the label for the first instance even if the budget T goes to infinity.*

According to Proposition 6, the deterministic KG is *not a consistent policy*, where the consistent policy refers to the policy that will provide correct labels for all instances (i.e., $H_T = H^*$) almost surely when T goes to infinity. We note that randomized KG policy can address this problem. However, from the proof of Proposition 6, randomized KG behaves similarly to the uniform sampling policy in many cases and its empirical performance is undesirable according to Section 8. In the next subsection, we will propose a new approximate allocation policy based on KG which is a consistent policy with superior empirical performance.

4.1 Optimistic Knowledge Gradient

The stage-wise reward can be viewed as a random variable with a two point distribution, i.e., with the probability $p_1 = \frac{a}{a+b}$ of being $R_1(a, b)$ and the probability $p_2 = \frac{b}{a+b}$ of being $R_2(a, b)$. The KG policy selects the instance with the largest *expected* reward. However, it is not consistent.

Algorithm 1 Optimistic Knowledge Gradient

Input: Parameters of prior distributions for instances $\{a_i^0, b_i^0\}_{i=1}^K$ and the budget T .

for $t = 0, \dots, T - 1$ **do**

Select the next instance i_t to label according to:

$$i_t = \arg \max_{i \in \{1, \dots, K\}} (R^+(a_i^t, b_i^t) \doteq \max(R_1(a_i^t, b_i^t), R_2(a_i^t, b_i^t))). \quad (17)$$

Acquire the label $y_{i_t} \in \{-1, 1\}$.

if $y_{i_t} = 1$ **then**

$$\text{padding-left: 4em;} a_{i_t}^{t+1} = a_{i_t}^t + 1, b_{i_t}^{t+1} = b_{i_t}^t; a_i^{t+1} = a_i^t, b_i^{t+1} = b_i^t \text{ for all } i \neq i_t.$$

else

$$\text{padding-left: 4em;} a_{i_t}^{t+1} = a_{i_t}^t, b_{i_t}^{t+1} = b_{i_t}^t + 1; a_i^{t+1} = a_i^t, b_i^{t+1} = b_i^t \text{ for all } i \neq i_t.$$

end if

end for

Output: The positive set $H_T = \{i : a_i^T \geq b_i^T\}$.




Figure 2: Illustration of $R^+(a, b)$.

In this section, we introduce a new index policy called “optimistic knowledge gradient” (Opt-KG) policy. The Opt-KG policy assumes that decision makers are optimistic in the sense that they select the next instance based on the optimistic outcome of the reward. As a simplest version of the Opt-KG policy, for any state (a_i^t, b_i^t) , the optimistic outcome of the reward $R^+(a_i^t, b_i^t)$ is defined as maximum over the reward of obtaining the label 1, $R_1(a_i^t, b_i^t)$, and the reward of obtaining the label -1 , $R_2(a_i^t, b_i^t)$. Then the optimistic decision maker selects the next instance i with the largest $R^+(a_i^t, b_i^t)$ as in (17) in Algorithm 1. The overall decision process using the Opt-KG policy is highlighted in Algorithm 1.

In the next theorem, we prove that Opt-KG policy is consistent.

Theorem 7 *Assuming that a_i^0 and b_i^0 are positive integers, the Opt-KG is a consistent policy, i.e., as T goes to infinity, the accuracy will be 100% (i.e., $H_T = H^*$) almost surely.*

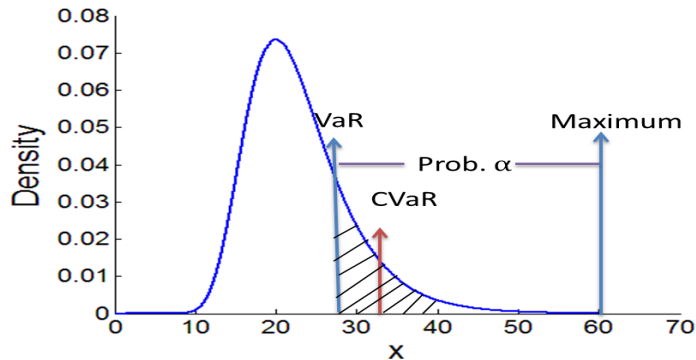


Figure 3: Illustration of Conditional Value-at-Risk.

The key of proving the consistency is to show that when T goes to infinity, each instance will be labeled infinitely many times. We prove this by showing that for any pair of positive integers (a, b) , $R^+(a, b) = \max(R_1(a, b), R_2(a, b)) > 0$ and $R^+(a, b) \rightarrow 0$ when $a + b \rightarrow \infty$. As an illustration, the values of $R^+(a, b)$ are plotted in Figure 2. Then, by strong law of large number, we obtain the consistency of the Opt-KG as stated in Theorem 7. The details are presented in the appendix. We have to note that asymptotic consistency is the minimum guarantee for a good policy. However, it does not necessarily guarantee the good empirical performance for the finite budget level. We will use experimental results to show the superior performance of the proposed policy.

The proposed Opt-KG policy is a general framework for budget allocation in crowd labeling. We can extend the allocation policy based on the maximum over the two possible rewards (Algorithm 1) to a more general policy using the conditional value-at-risk (CVaR) (Rockafellar and Uryasev, 2002). We note that here, instead of adopting the CVaR as a risk measure, we apply it to the reward distribution. In particular, for a random variable X with the support \mathcal{X} (e.g., the random reward with the two point distribution), let α -quantile function be denoted as $Q_\alpha(X) = \inf\{x \in \mathcal{X} : \alpha \leq F_X(x)\}$, where $F_X(\cdot)$ is the CDF of X . The value-at-risk $\text{VaR}_\alpha(X)$ is the smallest value such that the probability that X is less than (or equal to) it is greater than (or equal to) $1 - \alpha$: $\text{VaR}_\alpha(X) = Q_{1-\alpha}(X)$. The conditional value-at-risk ($\text{CVaR}_\alpha(X)$) is defined as the expected reward exceeding (or equal to) $\text{VaR}_\alpha(X)$. An illustration of CVaR is shown in Figure 3.

For our problem, according to Rockafellar and Uryasev (2002), $\text{CVaR}_\alpha(X)$ can be expressed as a simple linear program:

$$\begin{aligned} \text{CVaR}_\alpha(X) &= \max_{\{q_1 \geq 0, q_2 \geq 0\}} q_1 R_1 + q_2 R_2, \\ \text{s.t.} \quad q_1 &\leq \frac{1}{\alpha} p_1, \quad q_2 \leq \frac{1}{\alpha} p_2, \quad q_1 + q_2 = 1. \end{aligned}$$

As we can see, when $\alpha = 1$, $\text{CVaR}_\alpha(X) = p_1 R_1 + p_2 R_2$, which is the expected reward; when $\alpha \rightarrow 0$, $\text{CVaR}_\alpha(X) = \max(R_1, R_2)$, which is used as the selection criterion in (17) in Algorithm 1. In fact, a more general Opt-KG policy could be selecting the next instance with the largest $\text{CVaR}_\alpha(X)$ with a tuning parameter $\alpha \in [0, 1]$. We can extend Theorem 7 to prove that the policy based on $\text{CVaR}_\alpha(X)$ is consistent for any $\alpha < 1$. According to our

own experience, $\alpha \rightarrow 0$ usually has a better performance in our problem especially when the budget is very limited. Therefore, for the sake of presentation simplicity, we introduce the Opt-KG using $\max(R_1, R_2)$ (i.e., $\alpha \rightarrow 0$ in $\text{CVaR}_\alpha(X)$) as the selection criterion.

Finally, we highlight that the Opt-KG policy is computationally very efficient. For K instances with T units of the budget, the overall time and space complexity are $O(KT)$ and $O(K)$ respectively. It is much more efficient than the Gittins index policy which requires $O(T^6)$ time and space complexity.

4.2 Discussions

It is interesting to see the connection between the idea of making the decision based on the optimistic outcome of the reward and the UCB (upper confidence bounds) policy (Auer et al., 2002) for the classical multi-armed bandit problem as described in Section 3.3. In particular, the UCB policy selects the next arm with the maximum *upper confidence index*, which is defined as the current average reward plus the one-sided confidence interval. As we can see, the upper confidence index can be viewed as an “optimistic” estimate of the reward. However, we note that since we are in a Bayesian setting and our stage-wise rewards are artificially created and thus not *i.i.d.* for each arm, the UCB policy (Auer et al., 2002) cannot be directly applied to our problem.

In fact, our Opt-KG follows a more general principle of “optimism in the face uncertainty” (Szita and Lőrincz, 2008). Essentially, the non-consistency of KG is due to its nature of pure exploitation while a consistent policy should typically utilize exploration. One of the common techniques to handle the exploration-exploitation dilemma is to take an action based on an optimistic estimation of the rewards (see Szita and Lőrincz, 2008; Even-Dar and Mansour, 2001), which is the role $R^+(a, b)$ plays in Opt-KG.

For our problem, it is also straightforward to design the “pessimistic knowledge gradient” policy which selects the next instance i_t based on the pessimistic outcome of the reward, i.e., $i_t = \arg \max_i (R^-(a_i^t, b_i^t) \doteq \min(R_1(a_i^t, b_i^t), R_2(a_i^t, b_i^t)))$. However, as shown in the next proposition with the proof in the appendix, the pessimistic KG policy is inconsistent under the uniform prior.

Proposition 8 *When starting from the uniform prior (i.e., $a_i^0 = b_i^0 = 1$) for all θ_i , the pessimistic KG policy will acquire one label for each instance and then consistently acquire the label for the first instance even if the budget T goes to infinity.*

Finally, we discuss some other possible choices of prior distributions. For presentation simplicity, we only consider the Beta prior for each θ_i with the fixed parameters a_i^0 and b_i^0 . In practice, more complicated priors can be easily incorporated into our framework. For example, instead of using only one Beta prior, one can adopt a mixture of Beta distributions as the prior and the posterior will also follow a mixture of Beta distributions, which allows an easy inference about the posterior. As we show in the experiments (see Section 8.1.2), the uniform prior does not work well when the data is highly skewed in terms of class distribution. To address this problem, one possible choice is to adopt the prior $p(\theta) = w_1\text{Beta}(c, 1) + w_2\text{Beta}(1, 1) + w_3\text{Beta}(1, c)$ where w_1, w_2 and w_3 are the weights and c is a constant larger than 1 (e.g., $c = 5$). In such a prior, $B(c, 1)$ corresponds to the data with more positive labels while $B(1, c)$ to the data with more negative labels. In addition

to the mixture Beta prior, one can adopt the hierarchical Bayesian approach which puts hyper-priors on the parameters in the Beta priors. The inference can be performed using empirical Bayes approach (Gelman et al., 2013; Robert, 2007). In particular, one can periodically re-calculate the MAP estimate of the hyper-parameters based on the available data and update the model, but otherwise proceed with the given hyper-parameters. For common choices of hyper-priors of Beta, please refer to Section 5.3 in Gelman et al. (2013). These approaches can also be applied to model the workers’ reliability as we introduced in the next Section. For example, one can use a mixture of Beta distributions as the prior for the workers’ reliability, where $\text{Beta}(c, 1)$ corresponds to reliable workers, $\text{Beta}(1, 1)$ to random workers and $\text{Beta}(1, c)$ to malicious or poorly informed workers.

5. Incorporate Reliability of Heterogeneous Workers

In push crowdsourcing marketplaces, it is important to model workers’ reliability so that the decision maker could assign more instances to reliable workers. Assuming that there are M workers in a push marketplace, we can capture the reliability of the j -th worker by introducing an extra parameter $\rho_j \in [0, 1]$ as in (Dawid and Skene, 1979; Raykar et al., 2010; Karger et al., 2013b), which is defined as the probability of getting the same label as the one from a random fully reliable worker. Recall that the soft-label θ_i is the i -th instance’s probability of being labeled as positive by a fully reliable worker and let z_{ij} be the label provided by the j -th worker for the i -th instance. We model the distribution of z_{ij} for given θ_i and ρ_j using the *one-coin* model (Dawid and Skene, 1979; Karger et al., 2013b):

$$\begin{aligned} \Pr(z_{ij} = 1|\theta_i, \rho_j) &= \Pr(z_{ij} = 1|y_i = 1, \rho_j) \Pr(y_i = 1|\theta_i) + \Pr(z_{ij} = 1|y_i = -1, \rho_j) \Pr(y_i = -1|\theta_i) \\ &= \rho_j \theta_i + (1 - \rho_j)(1 - \theta_i); \end{aligned} \quad (18)$$

$$\begin{aligned} \Pr(z_{ij} = -1|\theta_i, \rho_j) &= \Pr(z_{ij} = -1|y_i = -1, \rho_j) \Pr(y_i = -1|\theta_i) + \Pr(z_{ij} = -1|y_i = 1, \rho_j) \Pr(y_i = 1|\theta_i) \\ &= \rho_j(1 - \theta_i) + (1 - \rho_j)\theta_i, \end{aligned} \quad (19)$$

where y_i denotes the label provided a random fully reliable worker for the i -th instance. We also note that it is straightforward to extend the current one-coin model to a more complex *two-coin* model (Dawid and Skene, 1979; Raykar et al., 2010) by introducing a pair of parameters (ρ_{j1}, ρ_{j2}) to model the j -th worker’s reliability. In particular, ρ_{j1} and ρ_{j2} are the probabilities of getting the positive and negative labels when a fully reliable worker provides the same label.

Here we make the following implicit assumption:

Assumption 9 *We assume that different workers make independent judgments and, for each single worker, the labels provided by him/her to different instances are also independent.*

As the parameter ρ_j increases from 0 to 1, the j -th worker’s reliability also increases in the sense that $\Pr(z_{ij} = 1|\theta_i, \rho_j)$ gets more and more close to θ_i , which is the probability of getting a positive label from a random fully reliable worker. Different types of workers can be easily characterized by ρ_j . When all $\rho_j = 1$, it recovers the previous model with fully reliable workers since $\Pr(z_{ij} = 1|\theta_i, \rho_j) = \theta_i$, i.e, each worker provides the label only according to the underlying soft-label of the instance. When $\rho_j = 0.5$, we have $\Pr(z_{ij} = 1|\theta_i, \rho_j) = \Pr(z_{ij} = -1|\theta_i, \rho_j) = 0.5$, which indicates that the j -th worker is a spammer, who randomly submits positive or negative labels. When $\rho_j = 0$, it indicates that the

j -th worker is poorly informed or misunderstands the instruction such that he/she always assigns wrong labels.

We assume that instances' soft-label $\{\theta_i\}_{i=1}^K$ and workers' reliability $\{\rho_j\}_{j=1}^M$ are drawn from known Beta prior distributions: $\theta_i \sim \text{Beta}(a_i^0, b_i^0)$ and $\rho_j \sim \text{Beta}(c_j^0, d_j^0)$. At each stage, we need to make the decision on both the next instance i to be labeled and the next worker j to label the instance i (we omit t in i, j here for notational simplicity). In other words, the action space $\mathcal{A} = \{(i, j) : (i, j) \in \{1, \dots, K\} \times \{1, \dots, M\}\}$. Once the decision is made, the distribution of the outcome z_{ij} is given by (18) and (19). Given the prior distributions and likelihood functions in (18) and (19), the Bayesian Markov Decision process can be formally defined as in Section 3. Similar to the homogeneous worker setting, the optimal inferred positive set H_T takes the form of $H_T = \{i : P_i^T \geq 0.5\}$ as in Proposition 2 with $P_i^t = \Pr(i \in H^* | \mathcal{F}_t) = \Pr(\theta_i \geq 0.5 | \mathcal{F}_t)$. The value function $V(S^0)$ still takes the form of (8), which can be further decomposed into the sum of stage-wise rewards in (9) using Proposition 4. Unfortunately, in the heterogeneous worker setting, the posterior distributions of θ_i and ρ_j are highly correlated with a sophisticated joint distribution, which makes the computation of stage-wise rewards in (9) much more challenging. In particular, given the prior $\theta_i \sim \text{Beta}(a_i^0, b_i^0)$ and $\rho_j \sim \text{Beta}(c_j^0, d_j^0)$, the posterior distribution of θ_i and ρ_j given the label $z_{ij} = z \in \{-1, 1\}$ takes the following form:

$$p(\theta_i, \rho_j | z_{ij} = z) = \frac{\Pr(z_{ij} = z | \theta_i, \rho_j) \text{Beta}(a_i^0, b_i^0) \text{Beta}(c_j^0, d_j^0)}{\Pr(z_{ij} = z)}, \quad (20)$$

where $\Pr(z_{ij} = z | \theta_i, \rho_j)$ is the likelihood function defined in (18) and (19) and

$$\begin{aligned} \Pr(z_{ij} = 1) &= \mathbb{E}(\Pr(z_{ij} = 1 | \theta_i, \rho_j)) = \mathbb{E}(\theta_i) \mathbb{E}(\rho_j) + (1 - \mathbb{E}(\theta_i))(1 - \mathbb{E}(\rho_j)) \\ &= \frac{a_i^0}{a_i^0 + b_i^0} \frac{c_j^0}{c_j^0 + d_j^0} + \frac{b_i^0}{a_i^0 + b_i^0} \frac{d_j^0}{c_j^0 + d_j^0}. \end{aligned}$$

As we can see, the posterior distribution $p(\theta_i, \rho_j | z_{ij} = z)$ no longer takes the form of the product of the distributions of θ_i and ρ_j and the marginal posterior of θ_i is no longer a Beta distribution. As a result, P_i^t does not have a simple representation as in (5), which makes the computation of the reward function much more difficult as the number of stages increases. Therefore, to apply our Opt-KG policy to large-scale applications, we need to use some approximate posterior inference techniques.

When applying Opt-KG, we need to perform $2 \cdot K \cdot M \cdot T$ inferences of the posterior distribution in total. Each approximate inference should be computed very efficiently, hopefully in a closed-form. For large-scale problems, most traditional approximate inference techniques such as Markov Chain Monte Carlo (MCMC) or variational Bayesian methods (e.g., Beal, 2003; Paisley et al., 2012) may lead to higher computational cost since each inference is an iterative procedure. To address the computational challenge, we apply the variational approximation with the moment matching technique so that each inference of the approximate posterior can be computed in a closed-form. In fact, any highly efficient approximate inference can be utilized to compute the reward function. Since the main focus of the paper is on the MDP model and Opt-KG policy, we omit the discussion for other possible approximate inference techniques. In particular, we first adopt the variational approximation by

Algorithm 2 Optimistic Knowledge Gradient for Heterogeneous Workers

Input: Parameters of prior distributions for instances $\{a_i^0, b_i^0\}_{i=1}^K$ and for workers $\{c_j^0, d_j^0\}_{j=1}^M$. The total budget T .

for $t = 0, \dots, T - 1$ **do**

1. Select the next instance i_t to label and the next worker j_t to label i_t according to:

$$(i_t, j_t) = \underset{(i,j) \in \{1, \dots, K\} \times \{1, \dots, M\}}{\arg \max} (R^+(a_i^t, b_i^t, c_j^t, d_j^t) \doteq \max(R_1(a_i^t, b_i^t, c_j^t, d_j^t), R_2(a_i^t, b_i^t, c_j^t, d_j^t))). \quad (21)$$

2. Acquire the label $z_{i_t j_t} \in \{-1, 1\}$ of the i -th instance from the j -th worker.

3. Update the posterior by setting:

$$a_{i_t}^{t+1} = \tilde{a}_{i_t}^t(z_{i_t j_t}) \quad b_{i_t}^{t+1} = \tilde{b}_{i_t}^t(z_{i_t j_t}) \quad c_{j_t}^{t+1} = \tilde{c}_{j_t}^t(z_{i_t j_t}) \quad d_{j_t}^{t+1} = \tilde{d}_{j_t}^t(z_{i_t j_t}),$$

and all parameters for $i \neq i_t$ and $j \neq j_t$ remain the same.

end for

Output: The positive set $H_T = \{i : a_i^T \geq b_i^T\}$.

assuming the conditional independence of θ_i and ρ_j :

$$p(\theta_i, \rho_j | z_{ij} = z) \approx p(\theta_i | z_{ij} = z) p(\rho_j | z_{ij} = z).$$

We further approximate $p(\theta_i | z_{ij} = z)$ and $p(\rho_j | z_{ij} = z)$ by two Beta distributions:

$$p(\theta_i | z_{ij} = z) \approx \text{Beta}(\tilde{a}_i(z), \tilde{b}_i(z)), \quad p(\rho_j | z_{ij} = z) \approx \text{Beta}(\tilde{c}_j(z), \tilde{d}_j(z)),$$

where the parameters $\tilde{a}_i(z)$, $\tilde{b}_i(z)$, $\tilde{c}_j(z)$, $\tilde{d}_j(z)$ are computed using moment matching with the analytical form presented in the appendix. After this approximation, the new posterior distributions of θ_i and ρ_j still have the same structure as their prior distribution, i.e., the product of two Beta distributions, which allows a repeatable use of this approximation every time when a new label is collected. Moreover, due to the Beta distribution approximation of $p(\theta_i | z_{ij} = z)$, the reward function takes a similar form as in the previous setting. In particular, assuming at a certain stage, θ_i has the posterior distribution $\text{Beta}(a_i, b_i)$ and ρ_j has the posterior distribution $\text{Beta}(c_j, d_j)$. The reward of getting positive and negative labels for the i -th instance from the j -th worker are presented in (22) and (23):

$$R_1(a_i, b_i, c_j, d_j) = h(I(\tilde{a}_i(z=1), \tilde{b}_i(z=1))) - h(I(a_i, b_i)), \quad (22)$$

$$R_2(a_i, b_i, c_j, d_j) = h(I(\tilde{a}_i(z=-1), \tilde{b}_i(z=-1))) - h(I(a_i, b_i)). \quad (23)$$

With the reward in place, we present Opt-KG for budget allocation in the heterogeneous worker setting in Algorithm 2. We also note that due to the variational approximation of the posterior, establishing the consistency results of Opt-KG becomes very challenging in the heterogeneous worker setting.

6. Extensions

Our MDP formulation is a general framework to address many complex settings of dynamic budget allocation problems in crowd labeling. In this section, we briefly discuss two

important extensions, where for both extensions, Opt-KG can be directly applied as an approximate policy. We note that for the sake of presentation simplicity, we only present these extensions in the noiseless homogeneous worker setting. Further extensions to the heterogeneous setting are rather straightforward using the technique from Section 5.

6.1 Utilizing Contextual Information

When the contextual information is available for instances, we could easily extend our model to incorporate such an important information. In particular, let the contextual information for the i -th instance be represented by a p -dimensional feature vector $\mathbf{x}_i \in \mathbb{R}^p$. We could utilize the feature information by assuming a logistic model for θ_i :

$$\theta_i \doteq \frac{\exp\{\langle \mathbf{w}, \mathbf{x}_i \rangle\}}{1 + \exp\{\langle \mathbf{w}, \mathbf{x}_i \rangle\}},$$

where \mathbf{w} is assumed to be drawn from a Gaussian prior $N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$. At the t -th stage with the current state $(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, the decision maker determines the instance i_t and acquire its label $y_{i_t} \in \{-1, 1\}$. Then we update the posterior $\boldsymbol{\mu}_{t+1}$ and $\boldsymbol{\Sigma}_{t+1}$ using the Laplace method as in Bayesian logistic regression (Bishop, 2007). Variational methods can be applied to further accelerate the posterior update (Jaakkola and Jordan, 2000). The details are provided in the appendix.

6.2 Multi-Class Categorization

Our MDP formulation can also be extended to deal with multi-class categorization problems, where each instance is a multiple choice question with several possible options (i.e., classes). More formally, in a multi-class setting with C different classes, we assume that the i -th instance is associated with a probability vector $\boldsymbol{\theta}_i = (\theta_{i1}, \dots, \theta_{iC})$, where θ_{ic} is the probability that the i -th instance will be labeled as the class c by a random fully reliable worker and $\sum_{c=1}^C \theta_{ic} = 1$. We assume that $\boldsymbol{\theta}_i$ has a Dirichlet prior $\boldsymbol{\theta}_i \sim \text{Dir}(\boldsymbol{\alpha}_i^0)$ and the initial state S^0 is a $K \times C$ matrix with $\boldsymbol{\alpha}_i^0$ as its i -th row. At each stage t with the current state S^t , we determine the next instance i_t to be labeled and collect its label $y_{i_t} \in \{1, \dots, C\}$, which follows the categorical distribution: $p(y_{i_t}) = \prod_{c=1}^C \theta_{i_t c}^{I(y_{i_t}=c)}$. Since the Dirichlet is the conjugate prior of the categorical distribution, the next state induced by the posterior distribution is: $S_{i_t}^{t+1} = S_{i_t}^t + \boldsymbol{\delta}_{y_{i_t}}$ and $S_i^{t+1} = S_i^t$ for all $i \neq i_t$. Here $\boldsymbol{\delta}_c$ is a row vector with one at the c -th entry and zeros at all other entries. The transition probability is:

$$\Pr(y_{i_t} = c | S^t, i_t) = \mathbb{E}(\theta_{i_t c} | S^t) = \frac{\alpha_{i_t c}^t}{\sum_{c=1}^C \alpha_{i_t c}^t}.$$

We denote the true set of instances in class c by $H_c^* = \{i : \theta_{ic} \geq \theta_{ic'}, \forall c' \neq c\}$. By a similar argument as in Proposition 2, at the final stage T , the estimated set of instances belonging to class c is

$$H_c^T = \{i : P_{ic}^T \geq P_{ic'}^T, \forall c' \neq c\},$$

where $P_{ic}^t = \Pr(i \in H_c^* | \mathcal{F}_t) = \Pr(\theta_{ic} \geq \theta_{ic'}, \forall c' \neq c | S^t)$. We note that if the i -th instance belongs to more than one H_c^T , we only assign it to the one with the smallest

index c so that $\{H_c^T\}_{c=1}^C$ forms a partition of $\{1, \dots, K\}$. Let $\mathbf{P}_i^t = (P_{i1}^t, \dots, P_{iC}^t)$ and $h(\mathbf{P}_i^t) = \max_{1 \leq c \leq C} P_{ic}^t$. The expected reward takes the form of:

$$R(S^t, i_t) = \mathbb{E}(h(\mathbf{P}_{i_t}^{t+1}) - h(\mathbf{P}_{i_t}^t) | S^t, i_t).$$

With the reward function in place, we can formulate the problem into a MDP and use DP to obtain the optimal policy and Opt-KG to compute an approximate policy. The only computational challenge is how to calculate P_{ic}^t efficiently so that the reward can be evaluated. We present an efficient method in the appendix. We can further use Dirichlet distribution to model workers reliability as in Liu and Wang (2012). Using multi-class Bayesian logistic regression, we can also incorporate contextual information into the multi-class setting in a straightforward manner.

7. Related Works

Categorical crowd labeling is one of the most popular tasks in crowdsourcing since it requires less effort of the workers to provide categorical labels than other tasks such as language translations. Most work in categorical crowd labeling are solving a static problem, i.e., inferring true labels and workers' reliability based on a static labeled data set (Dawid and Skene, 1979; Raykar et al., 2010; Liu and Wang, 2012; Welinder et al., 2010; Whitehill et al., 2009; Zhou et al., 2012; Liu et al., 2012; Gao and Zhou, 2013). The first work that incorporates diversity of worker reliability is by Dawid and Skene (1979), which uses EM to perform the point estimation on both worker reliability and true class labels. Based on that, Raykar et al. (2010) extended (Dawid and Skene, 1979) by introducing Beta prior for workers' reliability and features of instances in the binary setting; and Liu and Wang (2012) further introduced Dirichlet prior for modeling workers' reliability in the multi-class setting. Our work utilizes the modeling techniques in these two static models as basic building blocks but extends to dynamic budget allocation settings.

In recent years, there are several works that have been devoted into online learning or budget allocation in crowdsourcing (Karger et al., 2013a,b; Bachrach et al., 2012; Ho et al., 2013; Ertekin et al., 2012; Yan et al., 2011; Kamar et al., 2012; Ipeirotis et al., 2013). The method proposed in Karger et al. (2013b) is based on the one-coin model. In particular, it assigns instances to workers according to a random regular bipartite graph. Although the error rate is proved to achieve the minimax rate, its analysis is asymptotic and method is not optimal when the budget is limited. Karger et al. (2013a) further extended the work by Karger et al. (2013b) to the multi-class setting. The new labeling uncertainty method in Ipeirotis et al. (2013) is one of the state-of-the-art methods for repeated labeling. However, it does not model each worker's reliability and incorporate it into the allocation process. Ho et al. (2013) proposed an online primal dual method for adaptive task assignment and investigated the sample complexity to guarantee that the probability of making an error for each instance is less than a threshold. However, it requires gold samples to estimate workers' reliability. Kamar et al. (2012) used MDP to address a different decision problem in crowd labeling, where the decision maker collects labels for each instance one after another and only decides whether to hire an additional worker or not. Basically, it is an optimal stopping problem since there is no pre-fixed amount of budget and one needs to balance the accuracy v.s. the amount of budget. Since the accuracy and the amount of budget are in different

metrics, such a balance could be very subjective. Furthermore, the MDP framework in Kamar et al. (2012) cannot distinguish different workers. To the best of our knowledge, there is no existing method that characterizes the *optimal* allocation policy for finite T . In this work, with the MDP formulation and DP algorithm, we characterize the optimal policy for budget allocation in crowd labeling under any budget level.

We also note that the budget allocation in crowd labeling is fundamentally different from noisy active learning (Settles, 2009; Nowak, 2009). Active learning usually does not model the variability of labeling difficulties among instances and assumes a single (noisy) oracle; while in crowd labeling, we need to model both instances’ labeling difficulty and different workers’ reliability. Secondly, active learning requires the feature information of instances for the decision, which could be unavailable in crowd labeling. Finally, the goal of the active learning is to label as few instances as possible to learn a good classifier. In contrast, for budget allocation in crowd labeling, the goal is to infer the true labels for as many instances as possible.

In fact, our MDP formulation is essentially a finite-horizon Bayesian multi-armed bandit (MAB) problem. While the infinite-horizon Bayesian MAB has been well-studied and the optimal policy can be computed via Gittins index (Gittins, 1989), for finite-horizon Bayesian MAB, the Gittins index rule is only an approximate policy with high computational cost. The proposed Opt-KG and a more general conditional value-at-risk based KG could be general policies for Bayesian MAB. Recently, a Bayesian UCB policy was proposed to address a different Bayesian MAB problem (Kaufmann et al., 2012). However, it is not clear how to directly apply the policy to our problem since we are not updating the posterior of the mean of rewards as in Kaufmann et al. (2012). We note that our problem is also related to optimal stopping problem. The main difference is that the optimal stopping problem is infinite-horizon while our problem is finite-horizon and the decision process must stop when the budget is exhausted.

8. Experiments

In this section, we conduct empirical study to show some interesting properties of the proposed Opt-KG policy and compare its performance to other methods. We observe that several commonly used priors such as the uniform prior (Beta(1, 1)), Jeffery prior (Beta(1/2, 1/2)) and Haldane prior (Beta(0, 0)) for instances’ soft-label $\{\theta_i\}_{i=1}^K$ lead to very similar performance. Therefore, we adopt the uniform prior (Beta(1, 1)) unless otherwise specified. In addition, for the ease of comparison, the accuracy is defined as $\frac{|H_T \cap H^*| + |(H_T)^c \cap (H^*)^c|}{K}$, which is normalized between $[0, 1]$.

8.1 Simulation Study

In this section, we conduct simulated study. For each simulated experiment, we randomly generate 20 different sets of data and report the averaged accuracy. The deviations for different methods are similar and quite small and thus omitted for the purpose of better visualization and space-saving.




Figure 4: Labeling counts for instances with different levels of ambiguity.




Figure 5: Labeling counts for workers with different levels of reliability.

8.1.1 STUDY ON LABELING FREQUENCY

We first investigate that, in the homogeneous noiseless worker setting (i.e., workers are fully reliable), how the total budget is allocated among instances with different levels of ambiguity. In particular, we assume there are $K = 21$ instances with soft-labels $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \dots, \theta_K) = (0, 0.05, 0.1, \dots, 1)$. We vary the total budget $T = 5K, 15K, 50K$ and report the number of times that each instance is labeled on average over 20 independent runs. The results are presented in Figure 4. It can be seen from Figure 4 that, more ambiguous instances with θ close to 0.5 in general receive more labels than those simple instances with θ close to 0 or 1. A more interesting observation is that when the budget level is low (e.g., $T = 5K$ in Figure 4(a)), the policy spends less budget on those very ambiguous instances (e.g., $\theta = 0.45$ or 0.5), but more budget on exploring less ambiguous instances (e.g., $\theta = 0.35, 0.4$ or 0.6). When the budget goes higher (e.g., $T = 15K$ in Figure 4(b)), those very ambiguous instances receive more labels but the most ambiguous instance ($\theta = 0.5$) not necessarily receives the most labels. In fact, the instances with $\theta = 0.45$ and $\theta = 0.55$ receive more labels than that of the most ambiguous instance. When the total budget is sufficiently large (e.g., $T = 50K$ in Figure 4(c)), the most ambiguous instance receives the most labels since all the other instances have received enough labels to infer their true labels.




Figure 6: Density plot for different Beta distributions for generating each θ_i . Here, (a) represents there are more easier instances; (b) more ambiguous instances; (c) & (d) imbalanced class distributions with more positive instances.

Next, we investigate that, in the heterogeneous worker setting, how many instances each worker is assigned. We simulate $K = 21$ instances’ soft-labels as before and further simulate workers’ reliability $\boldsymbol{\rho} = (\rho_1, \rho_2, \dots, \rho_M) = (0.1, 0.15, \dots, 0.5, 0.505, 0.515, \dots, 0.995)$ for $M = 59$ workers. Such a simulation ensures that there are more reliable workers, which is in line with actual situation. We vary the total budget $T = 5K, 15K, 50K$ and report the number of instances that each worker is assigned on average over 20 independent runs in Figure 5. As one can see, when the budget level goes up, there is clear trend that more reliable workers receive more instances.

8.1.2 PRIOR FOR INSTANCES

We investigate how robust Opt-KG is when using the uniform prior for each θ_i . We first simulate $K = 50$ instances with each $\theta_i \sim \text{Beta}(0.5, 0.5)$, $\theta_i \sim \text{Beta}(2, 2)$, $\theta_i \sim \text{Beta}(2, 1)$ or $\theta_i \sim \text{Beta}(4, 1)$. The density functions of these four different Beta distributions are plotted in Figure 6. For each generating distribution of θ_i , we compare Opt-KG using the uniform prior ($\text{Beta}(1, 1)$) (in red line) to Opt-KG with the true generating distribution as the prior (in blue line). The comparison in accuracy with different levels of budget ($T = 2K, \dots, 20K$) is shown in Figure 7. As we can see, the performance of Opt-KG using two different priors are quite similar for most generating distributions except for $\theta_i \sim \text{Beta}(4, 1)$ (i.e., the highly imbalanced class distribution). When $\theta_i \sim \text{Beta}(4, 1)$, the Opt-KG with uniform prior needs at least $T = 16K$ units of budget to match the performance of Opt-KG with true generating distribution as the prior. This result indicates that for balanced class distributions, the uniform prior is a good choice and robust to the underlying distribution of θ_i . For highly imbalanced class distributions, if a uniform prior is adopted, one needs more budget to recover from the inaccurate prior belief.

8.1.3 PRIOR ON WORKERS

We investigate how sensitive the prior for the workers’ reliability ρ_j is. In particular, we simulate $K = 50$ instances with each $\theta_i \sim \text{Beta}(1, 1)$ and $M = 100$ workers with $\rho_j \sim \text{Beta}(3, 1)$, $\rho_j \sim \text{Beta}(8, 1)$ or $\rho_j \sim \text{Beta}(5, 2)$. We ensure that there are more reliable workers than spammers or poorly informed workers, which is in line with the actual situation. We use the




Figure 7: Comparison between Opt-KG using the uniform distribution and true generating distribution as the prior.




Figure 8: Density plot for different Beta distributions for generating ρ_j . The plot in (d) is the one that we use as the prior.

prior Beta(4, 1), which indicates that we have the prior belief that most workers perform reasonably well and the averaged accuracy is $4/5 = 80\%$. In Figure 8, we show different density functions for generating ρ_j and the prior that we use (in Figure 8 (d)). For each




Figure 9: Comparison between Opt-KG using Beta(4, 1) and true generating distribution prior as the prior.

generating distribution of θ_i , we compare the Opt-KG policy using the prior (Beta(4, 1)) (in red line) to the Opt-KG with the true generating distribution as the prior (in blue line). The comparison in accuracy with different levels of budget ($T = 2K, \dots, 20K$) is shown in Figure 9. From Figure 9, we observe that the performance of Opt-KG using two different priors are quite similar in all different settings. Hence, we will use Beta(4, 1) as the prior when the true prior of workers is unavailable.

8.1.4 PERFORMANCE COMPARISON UNDER THE HOMOGENEOUS NOISELESS WORKER SETTING

We compare the performance of Opt-KG under the homogeneous noiseless worker setting to several other competitors, including

1. Uniform: Uniform sampling.
2. KG(Random): Randomized knowledge gradient (Frazier et al., 2008).
3. Gittins-Inf: A Gittins-indexed based policy proposed by Xie and Frazier (2013) for solving an infinite-horizon Bayesian MAB problem where the reward is discounted by δ . Although it solves a different problem, we apply it as a heuristic by choosing the discount factor δ such that $T = 1/(1 - \delta)$.
4. NLU: The “new labeling uncertainty” method proposed by Ipeirotis et al. (2013).

We note that we do not compare to the finite-horizon Gittins index rule (Nino-Mora, 2011) since its computation is very expensive. On some small-scale problems, we observe that the finite-horizon Gittins index rule (Nino-Mora, 2011) has the similar performance as Gittins-Inf in Xie and Frazier (2013).

We simulate $K = 50$ instances with each $\theta_i \sim \text{Beta}(1, 1)$, $\theta_i \sim \text{Beta}(0.5, 0.5)$, $\theta_i \sim \text{Beta}(2, 2)$, $\theta_i \sim \text{Beta}(2, 1)$ or $\theta_i \sim \text{Beta}(4, 1)$ (see Figure 6). For each of the five settings, we vary the total budget $T = 2K, 3K, \dots, 20K$ and report the mean of accuracy for 20 independently generated sets of $\{\theta_i\}_{i=1}^K$. For the last four settings, we report the comparison among different methods when either using the uniform prior (“uni prior” for short) or the

true generating distribution as the prior. From Figure 10, the proposed Opt-KG outperforms all the other competitors in most settings regardless the choice of the prior. For $\theta_i \sim \text{Beta}(0.5, 0.5)$, NLU matches the performance of Opt-KG; and for $\theta_i \sim \text{Beta}(2, 2)$, Gittins-inf matches the performance of Opt-KG. We also observe that the performance of randomized KG only slightly improves that of uniform sampling.

8.1.5 PERFORMANCE COMPARISON UNDER THE HETEROGENEOUS WORKER SETTING

We compare the proposed Opt-KG under the heterogeneous worker setting to several other competitors:


1. Uniform: Uniform sampling.
2. KG(Random): Randomized knowledge gradient (Frazier et al., 2008).
3. KOS: The randomized budget allocation algorithm in Karger et al. (2013b).

We note that several competitors for the homogeneous worker setting (e.g., Gittins-inf and NLU) cannot be directly applied to the heterogeneous worker setting since they fail to model each worker’s reliability.


We simulate $K = 50$ instances with each $\theta_i \sim \text{Beta}(1, 1)$ and $M = 100$ workers with $\rho_j \sim \text{Beta}(4, 1)$, $\rho_j \sim \text{Beta}(3, 1)$, $\rho_j \sim \text{Beta}(8, 1)$ or $\rho_j \sim \text{Beta}(5, 2)$ (see Figure 8). For each of the four settings, we vary the total budget $T = 2K, 3K, \dots, 20K$ and report the mean of accuracy for 20 independently generated sets of parameters. For the last three settings, we report the comparison among different methods when either using $\text{Beta}(4, 1)$ prior or the true generating distribution for ρ_j as the prior. From Figure 11, the proposed Opt-KG outperforms all the other competitors regardless the choice of the prior.

8.2 Real Data


We compare different policies on a standard real data set for recognizing textual entailment (RTE) (Section 4.3 in Snow et al., 2008). There are 800 instances and each instance is a sentence pair. Each sentence pair is presented to 10 different workers to acquire binary choices of whether the second hypothesis sentence can be inferred from the first one. There are in total 164 different workers. We first consider the homogeneous noiseless setting without incorporating the diversity of workers and use the uniform prior ($\text{Beta}(1, 1)$) for each θ_i . In such a setting, once we decide to label an instance, we randomly choose a worker (who provides the label in the full data set) to acquire the label. Due to this randomness, we run each policy 20 times and report the mean of the accuracy in Figure 12(a). As we can see, Opt-KG, Gittins-inf and NLU all perform quite well. We also note that although Gittins-inf performs slightly better than our method on this data, it requires solving a linear system with $O(T^2)$ variables at each stage, which could be too expensive for large-scale applications. While our Opt-KG policy has a time complexity linear in KT and space complexity linear in K , which is much more efficient when a quick online decision is required. In particular, we present the comparison between Opt-KG and Gittins-inf on the averaged CPU time under different budget levels in Table 4. As one can see, Gittins-inf is computationally more expensive than Opt-KG.




(a) $\theta_i \sim \text{Beta}(1, 1)$ (True Prior) (b) $\theta_i \sim \text{Beta}(0.5, 0.5)$ (True Prior) (c) $\theta_i \sim \text{Beta}(0.5, 0.5)$ (Uni Prior)




(d) $\theta_i \sim \text{Beta}(2, 2)$ (True Prior)




(e) $\theta_i \sim \text{Beta}(2, 2)$ (Uni Prior)




(f) $\theta_i \sim \text{Beta}(2, 1)$ (True Prior)



(g) $\theta_i \sim \text{Beta}(2, 1)$ (Uni Prior)



(h) $\theta_i \sim \text{Beta}(4, 1)$ (True Prior)



(i) $\theta_i \sim \text{Beta}(4, 1)$ (Uni Prior)

Figure 10: Performance comparison under the homogeneous noiseless worker setting.

When the worker reliability is incorporated, we compare different policies in Figure 12(b). We put a $\text{Beta}(4, 1)$ prior distribution for each ρ_j which indicates that we have the

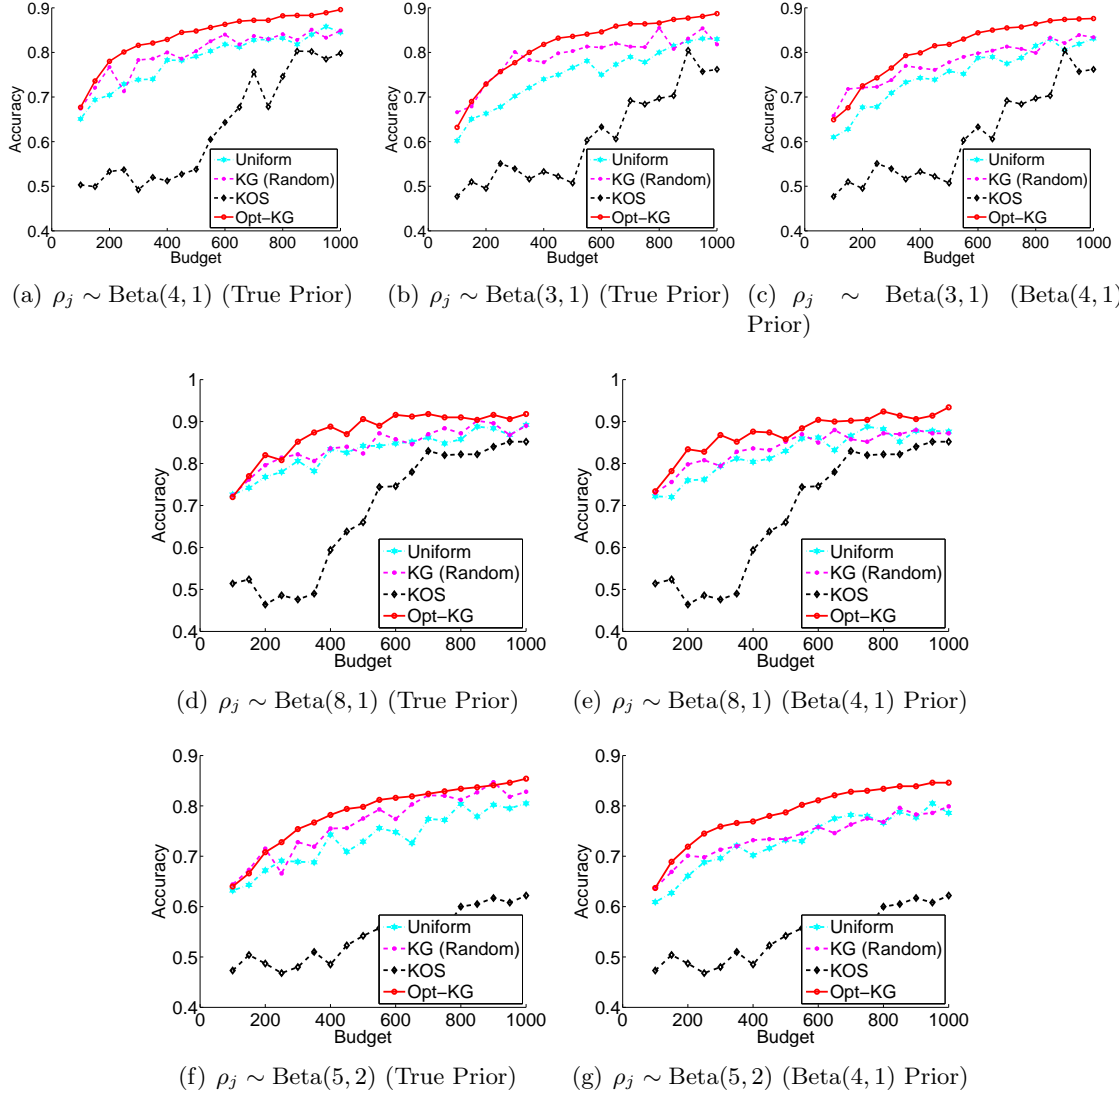


Figure 11: Performance comparison under the heterogeneous worker setting.

Budget T	$2K = 1,600$	$4K = 3,200$	$6K = 4,800$	$10K = 8,000$
Opt-KG	1.09	2.19	3.29	5.48
Gittins-inf	25.87	35.70	45.59	130.68

Table 4: Comparison in CPU time (seconds)

prior belief that most workers perform reasonably well. Other priors in Figure 8 lead to similar results and thus omitted here. As one can see, the accuracy of Opt-KG is much higher than that of other policies when T is small. It achieves the highest accuracy of 92.05% only using 40% of the total budget (i.e., on average, each instance is labeled 4 times). One may also observe that when $T > 4K = 3,200$, the performance of Opt-KG does not improve

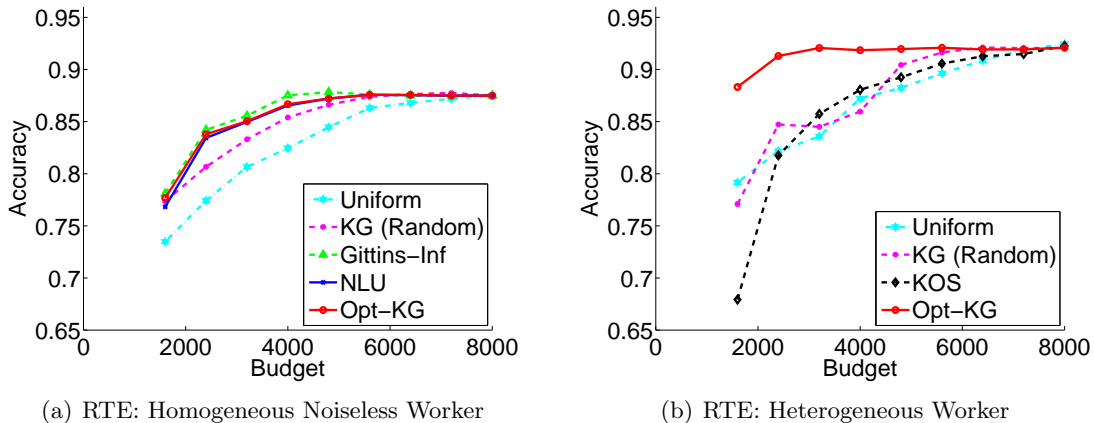


Figure 12: Performance comparison on the real data set.

and in fact, slightly downgrades a little bit. This is mainly due to the restrictiveness of the experimental setting. In particular, since the experiment is conducted on a fixed data set with partially observed labels, the Opt-KG cannot freely choose instance-worker pairs especially when the budget goes up (i.e., the action set is greatly restricted). According to our experience, such a phenomenon will not happen on experiments when labels can be obtained from any instance-worker pair. Comparing Figure 12(b) to 12(a), we also observe that Opt-KG under the heterogeneous worker setting performs much better than Opt-KG under the homogeneous worker setting, which indicates that it is beneficial to incorporate workers' reliability.

9. Conclusions and Future Work

In this paper, we propose to address the problem of budget allocation in crowd labeling. We model the problem using the Bayesian Markov decision process and characterize the optimal policy using the dynamic programming. We further propose a computationally more attractive approximate policy: optimistic knowledge gradient. Our MDP formulation is a general framework, which can be applied to binary or multi-class, contextual or non-contextual crowd labeling problems in either pull or push crowdsourcing marketplaces.

There are several possible future directions for this work. First, it is of great interest to show the consistency of Opt-KG in heterogeneous worker setting and further provide the theoretical results on the performance of Opt-KG under finite budget. Second, in this work, we assume that both instances and workers are equally priced. Although this assumption is standard in many crowd labeling applications, a dynamic pricing strategy as the allocation process proceeds will better motivate those more reliable workers to label more challenge instances. A recent work in Wang et al. (2013) provides some quality-based pricing algorithms for crowd workers and it will be interesting to incorporate their strategies into our dynamic allocation framework. Third, we assume that the labels provided by the same worker to different instances are independent. It is more interesting to consider that the workers' reliability will be improved during the labeling process when some useful feedback

can be provided. Further, since the proposed Opt-KG is a fairly general approximate policy for MDP, it is also interesting to apply it to other statistical decision problems.

Acknowledgments

We would like to thank Qiang Liu for sharing the code for KOS method; Jing Xie and Peter Frazier for sharing their code for computing infinite-horizon Gittins index; John Platt, Chris Burges and Kevin Murphy for helpful discussions; and anonymous reviewers and the associate editor for their constructive comments on improving the quality of the paper.

Appendix A. Proof of Results

In this section, we provide the proofs of the main results of our paper.

A.1 Proof of Proposition 2

The final positive set H_T is chosen to maximize the expected accuracy conditioned on \mathcal{F}_T :

$$H_T = \arg \max_H \mathbb{E} \left(\sum_{i=1}^K (\mathbf{1}(i \in H) \mathbf{1}(i \in H^*) + \mathbf{1}(i \notin H) \mathbf{1}(i \notin H^*)) \middle| \mathcal{F}_T \right). \quad (24)$$

According to the definition (6) of P_i^T , we can re-write (24) using the linearity of the expectation:

$$\begin{aligned} & \sum_{i=1}^K (\mathbf{1}(i \in H) \Pr(i \in H^* | \mathcal{F}_T) + \mathbf{1}(i \notin H) \Pr(i \notin H^* | \mathcal{F}_T)) \\ = & \sum_{i=1}^K (\mathbf{1}(i \in H) P_i^T + \mathbf{1}(i \notin H) (1 - P_i^T)). \end{aligned} \quad (25)$$

To maximize (25) over H , it is easy to see that we should set $i \in H$ if and only if $P_i^T \geq 0.5$. Therefore, we have the positive set

$$H_T = \{i : P_i^T \geq 0.5\}.$$

A.2 Proof of Corollary 3

Recall that

$$I(a, b) = \Pr(\theta \geq 0.5 | \theta \sim \text{Beta}(a, b)) = \frac{1}{B(a, b)} \int_{0.5}^1 t^{a-1} (1-t)^{b-1} dt, \quad (26)$$

where $B(a, b)$ is the beta function.

It is easy to see that $I(a, b) > 0.5 \iff I(a, b) > 1 - I(a, b)$. We re-write $1 - I(a, b)$ as follows

$$1 - I(a, b) = \frac{1}{B(a, b)} \int_0^{0.5} t^{a-1} (1-t)^{b-1} dt = \frac{1}{B(a, b)} \int_{0.5}^1 t^{b-1} (1-t)^{a-1} dt,$$

where the second equality is obtained by setting $t := 1 - t$. Then we have:

$$\begin{aligned} I(a, b) - (1 - I(a, b)) &= \frac{1}{B(a, b)} \int_{0.5}^1 (t^{a-1}(1-t)^{b-1} - t^{b-1}(1-t)^{a-1}) dt \\ &= \frac{1}{B(a, b)} \int_{0.5}^1 t^{a-1}(1-t)^{b-1} \left(\left(\frac{t}{1-t} \right)^{a-b} - 1 \right) dt. \end{aligned}$$

Since $t > 0.5$, $\frac{t}{1-t} > 1$. When $a > b$, $\left(\frac{t}{1-t}\right)^{a-b} > 1$ and hence $I(a, b) - (1 - I(a, b)) > 0$, i.e., $I(a, b) > 0.5$. When $a = b$, $\left(\frac{t}{1-t}\right)^{a-b} \equiv 1$ and $I(a, b) = 0.5$. When $a < b$, $\left(\frac{t}{1-t}\right)^{a-b} < 1$ and $I(a, b) < 0.5$.

A.3 Proof of Proposition 4

We use the proof technique in Xie and Frazier (2013) to prove Proposition 4. According to (8), the value function takes the following form,

$$V(S^0) = \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{i=1}^K h(P_i^T) \right). \quad (27)$$

To decompose the final accuracy $\sum_{i=1}^K h(P_i^T)$ into the incremental reward at each stage, we define $G_0 = \sum_{i=1}^K h(P_i^0)$ and $G_{t+1} = \sum_{i=1}^K h(P_i^{t+1}) - \sum_{i=1}^K h(P_i^t)$. Then, $\sum_{i=1}^K h(P_i^T)$ can be decomposed as: $\sum_{i=1}^K h(P_i^T) \equiv G_0 + \sum_{t=0}^{T-1} G_{t+1}$. The value function can now be re-written as follows:

$$\begin{aligned} V(S^0) &= G_0(S^0) + \sup_{\pi} \sum_{t=0}^{T-1} \mathbb{E}^{\pi}(G_{t+1}) \\ &= G_0(S^0) + \sup_{\pi} \sum_{t=0}^{T-1} \mathbb{E}^{\pi}(\mathbb{E}(G_{t+1}|\mathcal{F}_t)) \\ &= G_0(S^0) + \sup_{\pi} \sum_{t=0}^{T-1} \mathbb{E}^{\pi}(\mathbb{E}(G_{t+1}|S^t, i_t)). \end{aligned}$$

Here, the first inequality is true because G_0 is determinant and independent of π ; the second inequality is due to the tower property of conditional expectation and the third one holds because G_{t+1} , which is a function of P_i^{t+1} and P_i^t , depends on \mathcal{F}_t only through S^t and i_t . We define incremental expected reward gained by labeling the i_t -th instance at the state S^t as follows:

$$\begin{aligned} R(S^t, i_t) &= \mathbb{E}(G_{t+1}|S^t, i_t) = \mathbb{E} \left(\sum_{i=1}^K h(P_i^{t+1}) - \sum_{i=1}^K h(P_i^t) | S^t, i_t \right) \\ &= \mathbb{E} (h(P_{i_t}^{t+1}) - h(P_{i_t}^t) | S^t, i_t). \end{aligned} \quad (28)$$

The last equation is due to the fact that only $P_{i_t}^t$ will be changed if the i_t -th instance is labeled next. With the expected reward function in place, the value function in (8) can be re-formulated as:

$$V(S^0) = G_0(\mathbf{s}) + \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{t=0}^{T-1} R(S^t, i_t) \middle| S^0 \right). \quad (29)$$

A.4 Proof of Proposition 6

To prove the failure of deterministic KG, we first show a key property for the expected reward function:

$$R(a, b) = \frac{a}{a+b} (h(I(a+1, b)) - h(I(a, b))) + \frac{b}{a+b} (h(I(a, b+1)) - h(I(a, b))). \quad (30)$$

Lemma 10 *When a, b are positive integers, if $a = b$, $R(a, b) = \frac{0.5^{2a}}{aB(a, a)}$ and if $a \neq b$, $R(a, b) = 0$.*

To prove lemma 10, we first present several basic properties for $B(a, b)$ and $I(a, b)$, which will be used in all the following theorems and proofs.

1. Properties for $B(a, b)$:

$$B(a, b) = B(b, a), \quad (31)$$

$$B(a+1, b) = \frac{a}{a+b} B(a, b), \quad (32)$$

$$B(a, b+1) = \frac{b}{a+b} B(a, b). \quad (33)$$

2. Properties for $B(a, b)$:

$$I(a, b) = 1 - I(b, a), \quad (34)$$

$$I(a+1, b) = I(a, b) + \frac{0.5^{a+b}}{aB(a, b)}, \quad (35)$$

$$I(a, b+1) = I(a, b) - \frac{0.5^{a+b}}{bB(a, b)}. \quad (36)$$

The properties for $I(a, b)$ are derived from the basic property of regularized incomplete beta function.¹

Proof [Proof of Lemma 10]

When $a = b$, by Corollary 3, we have $I(a+1, b) > 0.5$, $I(a, b) = 0.5$ and $I(a, b+1) < 0.5$. Therefore, the expected reward (30) takes the following form:

$$\begin{aligned} R(a, b) &= 0.5(I(a+1, a) - I(a, a)) + 0.5((1 - I(a, a+1)) - I(a, a)) \\ &= I(a+1, a) - I(a, a) = \frac{0.5^{2a}}{aB(a, a)}. \end{aligned}$$

1. <http://dlmf.nist.gov/8.17>

When $a > b$, since a, b are integers, we have $a \geq b+1$ and hence $I(a+1, b) > 0.5, I(a, b) > 0.5, I(a, b+1) \geq 0.5$ according to Corollary 3. The expected reward (30) now becomes:

$$\begin{aligned}
 R(a, b) &= \frac{a}{a+b}I(a+1, b) + \frac{b}{a+b}I(a, b+1) - I(a, b) \\
 &= \frac{a}{a+b} \frac{1}{B(a+1, b)} \int_{0.5}^1 t \cdot t^{a-1}(1-t)^{b-1} dt \\
 &\quad + \frac{b}{a+b} \frac{1}{B(a, b+1)} \int_{0.5}^1 t^{a-1}(1-t)(1-t)^{b-1} dt - I(a, b) \\
 &= \frac{1}{B(a, b)} \int_{0.5}^1 (t + (1-t)) \cdot t^{a-1}(1-t)^{b-1} dt - I(a, b) \\
 &= I(a, b) - I(a, b) = 0.
 \end{aligned}$$

Here we use (32) and (33) to show that $\frac{a}{a+b} \frac{1}{B(a+1, b)} = \frac{b}{a+b} \frac{1}{B(a, b+1)} = \frac{1}{B(a, b)}$.

When $a \leq b-1$, we can prove $R(a, b) = 0$ in a similar way. ■

With Lemma 10 in place, the proof for Proposition 6 is straightforward. Recall that the deterministic KG policy chooses the next instance according to

$$i_t = \arg \max_i R(S^t, i) = \arg \max_i R(a_i^t, b_i^t),$$

and breaks the tie by selecting the one with the smallest index. Since $R(a, b) > 0$ if and only if $a = b$, at the initial stage $t = 0$, $R(a_i^0, b_i^0) > 0$ for those instances $i \in \mathcal{E} = \{i : a_i^0 = b_i^0\}$. The policy will first select $i_0 \in \mathcal{E}$ with the largest $R(a_{i_0}^0, b_{i_0}^0)$. After obtaining the label y_{i_0} , either $a_{i_0}^0$ or $b_{i_0}^0$ will add one and hence $a_{i_0}^1 \neq b_{i_0}^1$ and $R(a_{i_0}^1, b_{i_0}^1) = 0$. The policy will select another instance $i_1 \in \mathcal{E}$ with the “current” largest expected reward and the expected reward for i_1 after obtaining the label y_{i_1} will then become zero. As a consequence, the KG policy will label each instance in \mathcal{E} for the first $|\mathcal{E}|$ stages and $R(a_i^{|\mathcal{E}|}, b_i^{|\mathcal{E}|}) = 0$ for all $i \in \{1, \dots, K\}$. Then the deterministic policy will break the tie selecting the first instance to label. From now on, for any $t \geq |\mathcal{E}|$, if $a_1^t \neq b_1^t$, then the expected reward $R(a_1^t, b_1^t) = 0$. Since the expected reward for other instances are all zero, the policy will still label the first instance. On the other hand, if $a_1^t = b_1^t$, and the first instance is the only one with the positive expected reward and the policy will label it. Thus Proposition 6 is proved.

Remark 11 *For randomized KG, after getting one label for each instance in \mathcal{E} for the first $|\mathcal{E}|$ stages, the expected reward for each instance has become zero. Then randomized KG will uniformly select one instance to label. At any stage $t \geq |\mathcal{E}|$, if there exists one instance i (at most one instance) with $a_i^t = b_i^t$, the KG policy will provide the next label for i ; otherwise, it will randomly select an instance to label.*

A.5 Proof of Theorem 7

To prove the consistency of the Opt-KG policy, we first show the exact values for $R_\alpha^+(a, b) = \max(R_1(a, b), R_2(a, b))$.

1. When $a \geq b + 1$:

$$R_1(a, b) = I(a + 1, b) - I(a, b) = \frac{0.5^{a+b}}{aB(a, b)} > 0;$$

$$R_2(a, b) = I(a, b + 1) - I(a, b) = -\frac{0.5^{a+b}}{bB(a, b)} < 0.$$

Therefore,

$$R^+(a, b) = R_1(a, b) = \frac{0.5^{a+b}}{aB(a, b)} > 0.$$

2. When $a = b$:

$$R_1(a, b) = I(a + 1, a) - I(a, a) = \frac{0.5^{2a}}{aB(a, a)};$$

$$R_2(a, b) = 1 - I(a, a + 1) - I(a, a) = \frac{0.5^{2a}}{aB(a, a)}.$$

Therefore, we have $R_1 = R_2$ and

$$R^+(a, b) = R_1(a, b) = R_2(a, b) = \frac{0.5^{2a}}{aB(a, a)} > 0.$$

3. When $b - 1 \geq a$:

$$R_1(a, b) = I(a, b) - I(a + 1, b) = -\frac{0.5^{a+b}}{aB(a, b)} < 0;$$

$$R_2(a, b) = I(a, b) - I(a, b + 1) = \frac{0.5^{a+b}}{bB(a, b)} > 0.$$

Therefore

$$R^+(a, b) = R_2(a, b) = \frac{0.5^{a+b}}{bB(a, b)} > 0.$$

We note that the values of $R^+(a, b)$ for different a, b are plotted in Figure 2 in main text.

As we can see $R^+(a, b) > 0$, for any positive integers (a, b) , we first prove in the following lemma that

$$\lim_{a+b \rightarrow \infty} R^+(a, b) = 0. \tag{37}$$

Lemma 12 *Properties for $R^+(a, b)$:*

1. $R(a, b)$ is symmetric, i.e., $R^+(a, b) = R^+(b, a)$.
2. $\lim_{a \rightarrow \infty} R^+(a, a) = 0$.
3. For any fixed $a \geq 1$, $R^+(a + k, a - k) = R^+(a - k, a + k)$ is monotonically decreasing in k for $k = 0, \dots, a - 1$.

4. When $a \geq b$, for any fixed b , $R^+(a, b)$ is monotonically decreasing in a . By the symmetry of $R^+(a, b)$, when $b \geq a$, for any fixed a , $R^+(a, b)$ is monotonically decreasing in b .

By the above four properties, we have $\lim_{(a+b) \rightarrow \infty} R^+(a, b) = 0$.

Proof [Proof of Lemma 12]

We first prove these four properties.

- Property 1: By the fact that $B(a, b) = B(b, a)$, the symmetry of $R^+(a, b)$ is straightforward.
- Property 2: For $a > 1$, $\frac{R^+(a, a)}{R^+(a-1, a-1)} = \frac{2a-1}{2a} < 1$ and hence $R^+(a, a)$ is monotonically decreasing in a . Moreover,

$$R^+(a, a) = R^+(1, 1) \prod_{i=2}^a \frac{2i-1}{2i} = R^+(1, 1) \prod_{i=2}^a \left(1 - \frac{1}{2i}\right) \leq R^+(1, 1) e^{-\sum_{i=2}^a \frac{1}{2i}}.$$

Since $\lim_{a \rightarrow \infty} \sum_{i=2}^a \frac{1}{2i} = \infty$ and $R^+(a, a) \geq 0$, $\lim_{a \rightarrow \infty} R^+(a, a) = 0$.

- Property 3: For any $k \geq 0$,

$$\begin{aligned} \frac{R^+(a + (k + 1), a - (k + 1))}{R^+(a + k, a - k)} &= \frac{(a + k)B(a + k, a - k)}{(a + k + 1)B(a + (k + 1), a - (k + 1))} \\ &= \frac{a - (k + 1)}{a + (k + 1)} < 1. \end{aligned}$$

- Property 4: When $a \geq b$, for any fixed b :

$$\frac{R^+(a + 1, b)}{R^+(a, b)} = \frac{aB(a, b)}{2(a + 1)B(a + 1, b)} = \frac{a(a + b)}{2a(a + 1)} < 1.$$

According to the third property, when $a + b$ is an even number, we have $R^+(a, b) < R^+(\frac{a+b}{2}, \frac{a+b}{2})$. According to the fourth property, when $a + b$ is an odd number and $a \geq b + 1$, we have $R^+(a, b) < R^+(a - 1, b) < R^+(\frac{a+b-1}{2}, \frac{a+b-1}{2})$; while when $a + b$ is an odd number and $a \leq b - 1$, we have $R^+(a, b) < R^+(a, b - 1) < R^+(\frac{a+b-1}{2}, \frac{a+b-1}{2})$. Therefore,

$$R^+(a, b) < R^+\left(\lfloor \frac{a+b}{2} \rfloor, \lfloor \frac{a+b}{2} \rfloor\right).$$

According to the second property such that $\lim_{a \rightarrow \infty} R^+(a, a) = 0$, we obtain (37). ■

Using Lemma 12, we first show that, in any sample path, the Opt-KG will label each instance infinitely many times as T goes to infinity. Let $\eta_i(T)$ be a random variable representing the number of times that the i -th instance has been labeled until the stage T using Opt-KG. Given a sample path ω , let $\mathcal{I}(\omega) = \{i : \lim_{T \rightarrow \infty} \eta_i(T)(\omega) < \infty\}$ be the set of instances that has been labeled only finite number of times as T goes to infinity in

this sample path. We need to prove that $\mathcal{I}(\omega)$ is an empty set for any ω . We prove it by contradiction. Assuming that $\mathcal{I}(\omega)$ is not empty, then after a certain stage \widehat{T} , instances in $\mathcal{I}(\omega)$ will never be labeled. By Lemma 12, for any $j \in \mathcal{I}^c$, $\lim_{T \rightarrow \infty} R^+(a_j^T(\omega), b_j^T(\omega)) = 0$. Therefore, there will exist $\bar{T} > \widehat{T}$ such that:

$$\max_{j \in \mathcal{I}^c} R^+(a_j^{\bar{T}}(\omega), b_j^{\bar{T}}(\omega)) < \max_{i \in \mathcal{I}} R^+(a_i^{\widehat{T}}(\omega), b_i^{\widehat{T}}(\omega)) = \max_{i \in \mathcal{I}} R^+(a_i^{\bar{T}}(\omega), b_i^{\bar{T}}(\omega)).$$

Then according to the Opt-KG policy, the next instance to be labeled must be in $\mathcal{I}(\omega)$, which leads to the contradiction. Therefore, $\mathcal{I}(\omega)$ will be an empty set for any ω .

Let Y_i^s be the random variable which takes the value 1 if the s -th label of the i -th instance is 1 and the value -1 if the s -th label is 0. It is easy to see that $\mathbb{E}(Y_i^s | \theta_i) = \Pr(Y_i^s = 1 | \theta_i) = \theta_i$. Hence, Y_i^s , $s = 1, 2, \dots$ are independent and identically distributed random variables. By the fact that $\lim_{T \rightarrow \infty} \eta_T(i) = \infty$ in all sample paths and using the strong law of large number, we conclude that, conditioning on θ_i , $i = 1, \dots, K$, the conditional probability of

$$\lim_{T \rightarrow \infty} \frac{a_i^T - b_i^T}{\eta_i(T)} = \lim_{T \rightarrow \infty} \frac{\sum_{s=1}^{\eta_i(T)} Y_i^s}{\eta_i(T)} = \mathbb{E}(Y_i^s | \theta_i) = 2\theta_i - 1$$

for all $i = 1, \dots, K$, is one. According to Proposition 2, we have $H_T = \{i : a_i^T \geq b_i^T\}$ and $H^* = \{i : \theta_i \geq 0.5\}$. The accuracy is $\text{Acc}(T) = \frac{1}{K} (|H_T \cap H^*| + |H_T^c \cap (H^*)^c|)$. We have:

$$\begin{aligned} & \Pr\left(\lim_{T \rightarrow \infty} \text{Acc}(T) = 1 | \{\theta_i\}_{i=1}^K\right) = \Pr\left(\lim_{T \rightarrow \infty} (|H_T \cap H^*| + |H_T^c \cap (H^*)^c|) = K | \{\theta_i\}_{i=1}^K\right) \\ & \geq \Pr\left(\lim_{T \rightarrow \infty} \frac{a_i^T - b_i^T}{\eta_i(T)} = 2\theta_i - 1, \forall i = 1, \dots, K | \{\theta_i\}_{i=1}^K\right) = 1, \end{aligned}$$

whenever $\theta_i \neq 0.5$ for all i . The last inequality is due to the fact that, as long as θ_i is not 0.5 in any i , any sample path that gives the event $\lim_{T \rightarrow \infty} \frac{a_i^T - b_i^T}{\eta_i(T)} = 2\theta_i - 1, \forall i = 1, \dots, K$ also gives the event $\lim_{T \rightarrow \infty} (a_i^T - b_i^T) = \text{sgn}(2\theta_i - 1)(+\infty)$, which further implies $\lim_{T \rightarrow \infty} (|H_T \cap H^*| + |H_T^c \cap (H^*)^c|) = K$.

Finally, we have:

$$\begin{aligned} \Pr\left(\lim_{T \rightarrow \infty} \text{Acc}(T) = 1\right) &= \mathbb{E}_{\{\theta_i\}_{i=1}^K} \left[\Pr\left(\lim_{T \rightarrow \infty} \text{Acc}(T) = 1 | \{\theta_i\}_{i=1}^K\right) \right] \\ &= \mathbb{E}_{\{\theta_i: \theta_i \neq 0.5\}_{i=1}^K} \left[\Pr\left(\lim_{T \rightarrow \infty} \text{Acc}(T) = 1 | \{\theta_i\}_{i=1}^K\right) \right] \\ &= \mathbb{E}_{\{\theta_i: \theta_i \neq 0.5\}_{i=1}^K} [1] = 1, \end{aligned}$$

where the second equality is because $\{\theta_i : \exists i, \theta_i = 0.5\}$ is a zero measure set.

A.6 Proof of Proposition 8

Recall that our random reward is a two-point distribution with the probability $p_1 = \frac{a}{a+b}$ of being $R_1(a, b) = h(I(a+1, b)) - h(I(a, b))$ and $p_2 = \frac{b}{a+b}$ of being $R_2(a, b) = h(I(a, b+1)) - h(I(a, b))$. The pessimistic KG selects the next instance which maximizes $R^-(a, b) =$

$\min(R_1(a, b), R_2(a, b))$. To show that the policy is inconsistent, we first compute the exact values for $R^-(a, b)$ for positive integers (a, b) .

Utilizing Corollary 3 and the basic properties of $I(a, b)$ in (34), (35), (36), we have:

1. When $a \geq b + 1$:

$$R_1(a, b) = I(a + 1, b) - I(a, b) = \frac{0.5^{a+b}}{aB(a, b)} > 0;$$

$$R_2(a, b) = I(a, b + 1) - I(a, b) = -\frac{0.5^{a+b}}{bB(a, b)} < 0.$$

Therefore,

$$R^-(a, b) = R_2(a, b) = -\frac{0.5^{a+b}}{bB(a, b)} < 0.$$

2. When $a = b$:

$$R_1(a, b) = I(a + 1, a) - I(a, a) = \frac{0.5^{2a}}{aB(a, a)};$$

$$R_2(a, b) = 1 - I(a, a + 1) - I(a, a) = \frac{0.5^{2a}}{aB(a, a)}.$$

Therefore, we have $x_1 = x_2$ and

$$R^-(a, b) = R_1(a, b) = R_2(a, b) = \frac{0.5^{2a}}{aB(a, a)} > 0.$$

3. When $b - 1 \geq a$:

$$R_1(a, b) = I(a, b) - I(a + 1, b) = -\frac{0.5^{a+b}}{aB(a, b)} < 0;$$

$$R_2(a, b) = I(a, b) - I(a, b + 1) = \frac{0.5^{a+b}}{bB(a, b)} > 0.$$


Therefore

$$R^-(a, b) = R_1(a, b) = -\frac{0.5^{a+b}}{aB(a, b)} < 0.$$

We summarize the properties of $R^-(a, b)$ in the next Lemma.

Lemma 13 *Properties for $R^-(a, b)$:*

1. $R^-(a, b) > 0$ if and only if $a = b$.
2. $R^-(a, b)$ is symmetric, i.e., $R^-(a, b) = R^-(b, a)$
3. When $a = b + 1$, then $R^-(a, b) = R^-(b + 1, b)$ is monotonically increasing in b . By the symmetry of $R^-(a, b)$, when $b = a + 1$, $R^-(a, b) = R^-(a, a + 1)$ is monotonically increasing in a .

Figure 13: Illustration of $R^-(a, b)$.

4. When $a \geq b + 1$, for any fixed b , $R^-(a, b)$ is monotonically increasing in a . By the symmetry of $R^-(a, b)$, when $b \geq a + 1$, for any fixed a , $R^-(a, b)$ is monotonically increasing in b .

For better visualization, we plot values of $R^-(a, b)$ for different a, b in Figure 13. All the properties in Lemma 13 can be seen clearly from Figure 13. The proof of these properties are based on simple algebra and thus omitted here.

From Lemma 13, we can conclude that for any positive integers a, b with $a + b \neq 3$:

$$R^-(1, 2) = R^-(2, 1) < R^-(a, b). \quad (38)$$

Recall that the pessimistic KG selects:

$$i_t = \arg \max_{i \in \{1, \dots, K\}} R^-(a_i^t, b_i^t).$$

When starting from the uniform prior with $a_i^0 = b_i^0 = 1$ for all $i \in \{1 \dots, K\}$, the corresponding $R^-(a_i^0, b_i^0) = R^-(1, 1) > 0$. After obtaining a label for any instance i , the Beta parameters for θ_i will become either $(2, 1)$ or $(1, 2)$ with $R^-(1, 2) = R^-(2, 1) < 0$. Therefore, for the first K stages, the pessimistic KG policy will acquire the label for each instance once. For any instance i , we have either $a_i^K = 2, b_i^K = 1$ or $a_i^K = 1, b_i^K = 2$ at the stage K . Then the pessimistic KG policy will select the first instance to label. According to (38), for any $t \geq K$, $R^-(a_1^t, b_1^t) > R^-(1, 2) = R^-(2, 1)$. Therefore, the pessimistic KG policy will consistently acquire the label for the first instance. Since the tie will only appear at the stage K , the randomized pessimistic KG will also consistently select a single instance to label after K stages.

Appendix B. Incorporate Reliability of Heterogeneous Workers

As we discussed in Section 5 in main text, we approximate the posterior so that at any stage for all i, j , θ_i and ρ_j will follow Beta distributions. In particular, assuming at the

current state $\theta_i \sim \text{Beta}(a_i, b_i)$ and $\rho_j \sim \text{Beta}(c_j, d_j)$, the posterior distribution conditioned on z_{ij} takes the following form:

$$p(\theta_i, \rho_j | z_{ij} = 1) = \frac{\Pr(z_{ij} = 1 | \theta_i, \rho_j) \text{Beta}(a_i, b_i) \text{Beta}(c_j, d_j)}{\Pr(z_{ij} = 1)},$$

$$p(\theta_i, \rho_j | z_{ij} = -1) = \frac{\Pr(z_{ij} = -1 | \theta_i, \rho_j) \text{Beta}(a_i, b_i) \text{Beta}(c_j, d_j)}{\Pr(z_{ij} = -1)},$$

where the likelihood $\Pr(z_{ij} = z | \theta_i, \rho_j)$ for $z = 1, -1$ is defined in (18) and (19), i.e.,

$$\begin{aligned} \Pr(z_{ij} = 1 | \theta_i, \rho_j) &= \theta_i \rho_j + (1 - \theta_i)(1 - \rho_j), \\ \Pr(z_{ij} = -1 | \theta_i, \rho_j) &= (1 - \theta_i) \rho_j + \theta_i(1 - \rho_j). \end{aligned}$$

Also,

$$\begin{aligned} \Pr(z_{ij} = 1) &= \mathbb{E}(\Pr(z_{ij} = 1 | \theta_i, \rho_j)) = \mathbb{E}(\theta_i) \mathbb{E}(\rho_j) + (1 - \mathbb{E}(\theta_i))(1 - \mathbb{E}(\rho_j)) \\ &= \frac{a_i}{a_i + b_i} \frac{c_j}{c_j + d_j} + \frac{b_i}{a_i + b_i} \frac{d_j}{c_j + d_j}, \end{aligned}$$

$$\begin{aligned} \Pr(z_{ij} = -1) &= \mathbb{E}(\Pr(z_{ij} = -1 | \theta_i, \rho_j)) = (1 - \mathbb{E}(\theta_i)) \mathbb{E}(\rho_j) + \mathbb{E}(\theta_i)(1 - \mathbb{E}(\rho_j)) \\ &= \frac{b_i}{a_i + b_i} \frac{c_j}{c_j + d_j} + \frac{a_i}{a_i + b_i} \frac{d_j}{c_j + d_j}. \end{aligned}$$

The posterior distributions $p(\theta_i, \rho_j | z_{ij} = z)$ no longer takes the form of the product of Beta distributions on θ_i and ρ_j . Therefore, we use variational approximation by first assuming the conditional independence of θ_i and ρ_j :

$$p(\theta_i, \rho_j | z_{ij} = z) \approx p(\theta_i | z_{ij} = z) p(\rho_j | z_{ij} = z).$$

In fact, the exact form of marginal distributions can be calculated as follows:

$$\begin{aligned} p(\theta_i | z_{ij} = 1) &= \frac{\theta_i \mathbb{E}(\rho_j) + (1 - \theta_i)(1 - \mathbb{E}(\rho_j))}{\Pr(z_{ij} = 1)} \text{Beta}(a_i, b_i), \\ p(\rho_j | z_{ij} = 1) &= \frac{\mathbb{E}(\theta_i) \rho_j + (1 - \mathbb{E}(\theta_i))(1 - \rho_j)}{\Pr(z_{ij} = 1)} \text{Beta}(c_j, d_j), \\ p(\theta_i | z_{ij} = -1) &= \frac{(1 - \theta_i) \mathbb{E}(\rho_j) + \theta_i(1 - \mathbb{E}(\rho_j))}{\Pr(z_{ij} = -1)} \text{Beta}(a_i, b_i), \\ p(\rho_j | z_{ij} = -1) &= \frac{(1 - \mathbb{E}(\theta_i)) \rho_j + \mathbb{E}(\theta_i)(1 - \rho_j)}{\Pr(z_{ij} = -1)} \text{Beta}(c_j, d_j). \end{aligned}$$

To approximate the marginal distribution as Beta distribution, we use the moment matching technique. In particular, we approximate $p(\theta_i | z_{ij} = z) \approx \text{Beta}(\tilde{a}_i(z), \tilde{b}_i(z))$ such that

$$\tilde{\mathbb{E}}_z(\theta_i) \doteq \mathbb{E}_{p(\theta_i | z_{ij} = z)}(\theta_i) = \frac{\tilde{a}_i(z)}{\tilde{a}_i(z) + \tilde{b}_i(z)}, \quad (39)$$

$$\tilde{\mathbb{E}}_z(\theta_i^2) \doteq \mathbb{E}_{p(\theta_i | z_{ij} = z)}(\theta_i^2) = \frac{\tilde{a}_i(z)(\tilde{a}_i(z) + 1)}{(\tilde{a}_i(z) + \tilde{b}_i(z))(\tilde{a}_i(z) + \tilde{b}_i(z) + 1)}, \quad (40)$$

where $\frac{\tilde{a}_i(z)}{\tilde{a}_i(z)+\tilde{b}_i(z)}$ and $\frac{\tilde{a}_i(z)(\tilde{a}_i(z)+1)}{(\tilde{a}_i(z)+\tilde{b}_i(z))(\tilde{a}_i(z)+\tilde{b}_i(z)+1)}$ are the first and second order moment of Beta($\tilde{a}_i(z), \tilde{b}_i(z)$). To make (39) and (40) hold, we have:

$$\tilde{a}_i(z) = \tilde{\mathbb{E}}_z(\theta_i) \frac{\tilde{\mathbb{E}}_z(\theta_i) - \tilde{\mathbb{E}}_z(\theta_i^2)}{\tilde{\mathbb{E}}_z(\theta_i^2) - \left(\tilde{\mathbb{E}}_z(\theta_i)\right)^2}, \quad (41)$$

$$\tilde{b}_i(z) = (1 - \tilde{\mathbb{E}}_z(\theta_i)) \frac{\tilde{\mathbb{E}}_z(\theta_i) - \tilde{\mathbb{E}}_z(\theta_i^2)}{\tilde{\mathbb{E}}_z(\theta_i^2) - \left(\tilde{\mathbb{E}}_z(\theta_i)\right)^2}. \quad (42)$$

Similarly, we approximate $p(\rho_j | z_{ij} = z) \approx \text{Beta}(\tilde{c}_j(z), \tilde{d}_j(z))$, such that

$$\tilde{\mathbb{E}}_z(\rho_j) \doteq \mathbb{E}_{p(\rho_j | z_{ij}=z)}(\rho_j) = \frac{\tilde{c}_j(z)}{\tilde{c}_j(z) + \tilde{d}_j(z)}, \quad (43)$$

$$\tilde{\mathbb{E}}_z(\rho_j^2) \doteq \mathbb{E}_{p(\rho_j | z_{ij}=z)}(\rho_j^2) = \frac{\tilde{c}_j(z)(\tilde{c}_j(z) + 1)}{(\tilde{c}_j(z) + \tilde{d}_j(z))(\tilde{c}_j(z) + \tilde{d}_j(z) + 1)}, \quad (44)$$

where $\frac{\tilde{c}_j(z)}{\tilde{c}_j(z)+\tilde{d}_j(z)}$ and $\frac{\tilde{c}_j(z)(\tilde{c}_j(z)+1)}{(\tilde{c}_j(z)+\tilde{d}_j(z))(\tilde{c}_j(z)+\tilde{d}_j(z)+1)}$ are the first and second order moment of Beta($\tilde{c}_j(z), \tilde{d}_j(z)$). To make (39) and (40) hold, we have:

$$\tilde{c}_j(z) = \tilde{\mathbb{E}}_z(\rho_j) \frac{\tilde{\mathbb{E}}_z(\rho_j) - \tilde{\mathbb{E}}_z(\rho_j^2)}{\tilde{\mathbb{E}}_z(\rho_j^2) - \left(\tilde{\mathbb{E}}_z(\rho_j)\right)^2}, \quad (45)$$

$$\tilde{d}_j(z) = (1 - \tilde{\mathbb{E}}_z(\rho_j)) \frac{\tilde{\mathbb{E}}_z(\rho_j) - \tilde{\mathbb{E}}_z(\rho_j^2)}{\tilde{\mathbb{E}}_z(\rho_j^2) - \left(\tilde{\mathbb{E}}_z(\rho_j)\right)^2}. \quad (46)$$

Furthermore, we can compute the exact values for $\tilde{\mathbb{E}}_z(\theta_i)$, $\tilde{\mathbb{E}}_z(\theta_i^2)$, $\tilde{\mathbb{E}}_z(\rho_j)$ and $\tilde{\mathbb{E}}_z(\rho_j^2)$ as follows.

$$\begin{aligned} \tilde{\mathbb{E}}_1(\theta_i) &= \frac{\mathbb{E}(\theta_i^2)\mathbb{E}(\rho_j) + (\mathbb{E}(\theta_i) - \mathbb{E}(\theta_i^2))(1 - \mathbb{E}(\rho_j))}{p(z_{ij} = 1)} = \frac{a_i((a_i + 1)c_j + b_i d_j)}{(a_i + b_i + 1)(a_i c_j + b_i d_j)}, \\ \tilde{\mathbb{E}}_1(\theta_i^2) &= \frac{\mathbb{E}(\theta_i^3)\mathbb{E}(\rho_j) + (\mathbb{E}(\theta_i^2) - \mathbb{E}(\theta_i^3))(1 - \mathbb{E}(\rho_j))}{p(z_{ij} = 1)} = \frac{a_i(a_i + 1)((a_i + 2)c_j + b_i d_j)}{(a_i + b_i + 1)(a_i + b_i + 2)(a_i c_j + b_i d_j)}, \\ \tilde{\mathbb{E}}_{-1}(\theta_i) &= \frac{(\mathbb{E}(\theta_i) - \mathbb{E}(\theta_i^2))\mathbb{E}(\rho_j) + \mathbb{E}(\theta_i^2)(1 - \mathbb{E}(\rho_j))}{p(z_{ij} = -1)} = \frac{a_i(b_i c_j + (a_i + 1)d_j)}{(a_i + b_i + 1)(b_i c_j + a_i d_j)}, \\ \tilde{\mathbb{E}}_{-1}(\theta_i^2) &= \frac{(\mathbb{E}(\theta_i^2) - \mathbb{E}(\theta_i^3))\mathbb{E}(\rho_j) + \mathbb{E}(\theta_i^3)(1 - \mathbb{E}(\rho_j))}{p(z_{ij} = -1)} = \frac{a_i(a_i + 1)(b_i c_j + (a_i + 2)d_j)}{(a_i + b_i + 1)(a_i + b_i + 2)(b_i c_j + a_i d_j)}, \\ \tilde{\mathbb{E}}_1(\rho_j) &= \frac{\mathbb{E}(\theta_i)\mathbb{E}(\rho_j^2) + (1 - \mathbb{E}(\theta_i))(\mathbb{E}(\rho_j) - \mathbb{E}(\rho_j^2))}{p(z_{ij} = 1)} = \frac{c_j(a_i(c_j + 1) + b_i d_j)}{(c_j + d_j + 1)(a_i c_j + b_i d_j)}, \\ \tilde{\mathbb{E}}_1(\rho_j^2) &= \frac{\mathbb{E}(\theta_i)\mathbb{E}(\rho_j^3) + (1 - \mathbb{E}(\theta_i))(\mathbb{E}(\rho_j^2) - \mathbb{E}(\rho_j^3))}{p(z_{ij} = 1)} = \frac{c_j(c_j + 1)(a_i(c_j + 2) + b_i d_j)}{(c_j + d_j + 1)(c_j + d_j + 2)(a_i c_j + b_i d_j)}, \\ \tilde{\mathbb{E}}_{-1}(\rho_j) &= \frac{(1 - \mathbb{E}(\theta_i))\mathbb{E}(\rho_j^2) + \mathbb{E}(\theta_i)(\mathbb{E}(\rho_j) - \mathbb{E}(\rho_j^2))}{p(z_{ij} = -1)} = \frac{c_j(b_i(c_j + 1) + a_i d_j)}{(c_j + d_j + 1)(b_i c_j + a_i d_j)}, \\ \tilde{\mathbb{E}}_{-1}(\rho_j^2) &= \frac{(1 - \mathbb{E}(\theta_i))\mathbb{E}(\rho_j^3) + \mathbb{E}(\theta_i)(\mathbb{E}(\rho_j^2) - \mathbb{E}(\rho_j^3))}{p(z_{ij} = -1)} = \frac{c_j(c_j + 1)(b_i(c_j + 2) + a_i d_j)}{(c_j + d_j + 1)(c_j + d_j + 2)(b_i c_j + a_i d_j)}. \end{aligned}$$

Assuming that at a certain stage, θ_i follows a Beta posterior $\text{Beta}(a_i, b_i)$ and ρ_j follows a Beta posterior $\text{Beta}(c_j, d_j)$, the reward of getting positive and negative labels for the i -th instance from the j -th worker are:

$$R_1(a_i, b_i, c_j, d_j) = h(I(\tilde{a}_i(z=1), \tilde{b}_i(z=1))) - h(I(a_i, b_i)), \quad (47)$$

$$R_2(a_i, b_i, c_j, d_j) = h(I(\tilde{a}_i(z=-1), \tilde{b}_i(z=-1))) - h(I(a_i, b_i)), \quad (48)$$

where $\tilde{a}_i(z = \pm 1)$ and $\tilde{b}_i(z = \pm 1)$ are defined in (41) and (42), which further depend on c_j and d_j through $\tilde{\mathbb{E}}_z(\theta_i)$ and $\tilde{\mathbb{E}}_z(\theta_i^2)$. With the reward in place, we can directly apply the Opt-KG policy in the heterogeneous worker setting.

Appendix C. Extensions

In this section, we provide detailed Opt-KG algorithms for extensions in Section 6.

C.1 Utilizing Contextual Information

When each instance is associated with a p -dimensional feature vector $\mathbf{x}_i \in \mathbb{R}^p$, we incorporate the feature information in our budget allocation problem by assuming:

$$\theta_i = \sigma(\langle \mathbf{w}, \mathbf{x}_i \rangle) \doteq \frac{1}{1 + \exp\{-\langle \mathbf{w}, \mathbf{x}_i \rangle\}}, \quad (49)$$

where $\sigma(x) = \frac{1}{1 + \exp\{-x\}}$ is the sigmoid function and \mathbf{w} is assumed to be drawn from a Gaussian prior $N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$. At the t -th stage with the state $S^t = (\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ and $\mathbf{w} \sim (\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$, the decision maker chooses the i_t -th instance to be labeled and observes the label $y_{i_t} \in \{-1, 1\}$. The posterior distribution $p(\mathbf{w}|y_{i_t}, S^t) \propto p(y_{i_t}|\mathbf{w})p(\mathbf{w}|S^t)$ has the following log-likelihood:

$$\begin{aligned} \ln p(\mathbf{w}|y_{i_t}, S^t) &= \ln p(y_{i_t}|\mathbf{w}) + \ln p(\mathbf{w}|S^t) + \text{const} \\ &= \mathbf{1}(y_{i_t} = 1) \ln \sigma(\langle \mathbf{w}, \mathbf{x}_{i_t} \rangle) + \mathbf{1}(y_{i_t} = -1) \ln (1 - \sigma(\langle \mathbf{w}, \mathbf{x}_{i_t} \rangle)) \\ &\quad - \frac{1}{2} (\mathbf{w} - \boldsymbol{\mu}_t)' \boldsymbol{\Omega}_t (\mathbf{w} - \boldsymbol{\mu}_t) + \text{const}, \end{aligned}$$

where $\boldsymbol{\Omega}_t = (\boldsymbol{\Sigma}_t)^{-1}$ is the precision matrix. To approximate $p(\mathbf{w}|y_{i_t}, \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)$ by a Gaussian distribution $N(\boldsymbol{\mu}_{t+1}, \boldsymbol{\Sigma}_{t+1})$, we use the Laplace method (see Chapter 4.4 in Bishop, 2007). In particular, the mean of the posterior Gaussian is the MAP (maximum a posteriori) estimator of \mathbf{w} :

$$\boldsymbol{\mu}_{t+1} = \arg \max_{\mathbf{w}} \ln p(\mathbf{w}|y_{i_t}, S^t), \quad (50)$$

which can be computed by any numerical optimization method (e.g., Newton's method). The precision matrix takes the following form,

$$\boldsymbol{\Omega}_{t+1} = -\nabla^2 \ln p(\mathbf{w}|y_{i_t}, S^t) \Big|_{\mathbf{w}=\boldsymbol{\mu}_{t+1}} = \boldsymbol{\Omega}_t + \sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_{t+1}}) (1 - \sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_{t+1}})) \mathbf{x}_{i_{t+1}} \mathbf{x}'_{i_{t+1}}.$$

By Sherman-Morrison formula, the covariance matrix can be computed as,

$$\boldsymbol{\Sigma}_{t+1} = (\boldsymbol{\Omega}_{t+1})^{-1} = \boldsymbol{\Sigma}_t - \frac{\sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_t}) (1 - \sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_t}))}{1 + \sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_t}) (1 - \sigma(\boldsymbol{\mu}'_{t+1} \mathbf{x}_{i_t})) \mathbf{x}'_{i_t} \boldsymbol{\Sigma}_t \mathbf{x}_{i_t}} \boldsymbol{\Sigma}_t \mathbf{x}_{i_t+1} \mathbf{x}'_{i_t} \boldsymbol{\Sigma}_t.$$

We also calculate the transition probability of $y_{i_t} = 1$ and $y_{i_t} = -1$ using the technique from Bayesian logistic regression (see Chapter 4.5 in Bishop, 2007):

$$\Pr(y_{i_t} = 1|S^t, i_t) = \int p(y_{i_t} = 1|\mathbf{w})p(\mathbf{w}|S^t)d\mathbf{w} = \int \sigma(\mathbf{w}'\mathbf{x}_i)p(\mathbf{w}|S^t)d\mathbf{w} \approx \sigma(\mu_i\kappa(s_i^2)),$$

where $\kappa(s_i^2) = (1 + \pi s_i^2/8)^{-1/2}$ and $\mu_i = \langle \boldsymbol{\mu}_t, \mathbf{x}_i \rangle$ and $s_i^2 = \mathbf{x}_i' \boldsymbol{\Sigma}_t \mathbf{x}_i$.

To calculate the reward function, in addition to the transition probability, we also need to compute:

$$\begin{aligned} P_i^t &= \Pr(\theta_i \geq 0.5|\mathcal{F}_t) \\ &= \Pr\left(\frac{1}{1 + \exp\{-\mathbf{w}'_t\mathbf{x}_i\}} \geq 0.5 \mid \mathbf{w}_t \sim N(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)\right) \\ &= \Pr(\mathbf{w}'_t\mathbf{x}_i \geq 0 \mid \mathbf{w}_t \sim N(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)) \\ &= \int_0^\infty \left(\int_{\mathbf{w}} \delta(c - \langle \mathbf{w}, \mathbf{x}_i \rangle) N(\mathbf{w}|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t) d\mathbf{w} \right) dc, \end{aligned}$$

where $\delta(\cdot)$ is the Dirac delta function. Let

$$p(c) = \int_{\mathbf{w}} \delta(c - \langle \mathbf{w}, \mathbf{x}_i \rangle) N(\mathbf{w}|\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t) d\mathbf{w}.$$

Since the marginal of a Gaussian distribution is still a Gaussian, $p(c)$ is a univariate-Gaussian distribution with the mean and variance:

$$\begin{aligned} \mu_i &= \mathbb{E}(c) = \langle \mathbb{E}(\mathbf{w}), \mathbf{x}_i \rangle = \langle \boldsymbol{\mu}_t, \mathbf{x}_i \rangle, \\ s_i^2 &= \text{Var}(c) = (\mathbf{x}_i)' \text{Cov}(\mathbf{w}, \mathbf{w}) \mathbf{x}_i = (\mathbf{x}_i)' \boldsymbol{\Sigma}_t \mathbf{x}_i. \end{aligned}$$

Therefore, we have:

$$P_i^t = \int_0^\infty p(c)dc = 1 - \Phi\left(-\frac{\mu_i}{s_i}\right), \quad (51)$$

where $\Phi(\cdot)$ is the CDF of the standard Gaussian distribution.

With P_i^t and transition probability in place, the expected reward in value function takes the following form :

$$R(S^t, i_t) = \mathbb{E}\left(\sum_{i=1}^K h(P_i^{t+1}) - \sum_{i=1}^K h(P_i^t) \mid S^t, i_t\right). \quad (52)$$

We note that since \mathbf{w} will affect all P_i^t , the summation from 1 to K in (52) can not be omitted and hence (52) cannot be written as $\mathbb{E}(h(P_{i_t}^{t+1}) - h(P_{i_t}^t) \mid S^t, i_t)$ in (28). In this problem, KG or Opt-KG need to solve $O(2TK)$ optimization problems to compute the mean of the posterior as in (50), which could be computationally quite expensive. One possibility to address this problem is to use the variational Bayesian logistic regression (Jaakkola and Jordan, 2000), which could lead to a faster optimization procedure.

C.2 Multi-Class Categorization

Given the model and notations introduced in Section 6.2, at the final stage T when all budget is used up, we construct the set H_c^T for each class c to maximize the conditional expected classification accuracy:

$$\begin{aligned} \{H_c^T\}_{c=1}^C &= \arg \max_{H_c \subseteq \{1, \dots, C\}, H_c \cap H_{\tilde{c}} = \emptyset} \mathbb{E} \left(\sum_{i=1}^K \sum_{c=1}^C I(i \in H_c) I(i \in H_c^*) \middle| \mathcal{F}_T \right) \\ &= \arg \max_{H_c \subseteq \{1, \dots, C\}, H_c \cap H_{\tilde{c}} = \emptyset} \sum_{i=1}^K \sum_{c=1}^C I(i \in H_c) \Pr(i \in H_c^* | \mathcal{F}_T). \end{aligned} \quad (53)$$

Here, $H_c^* = \{i : \theta_{ic} \geq \theta_{i\tilde{c}}, \forall \tilde{c} \neq c\}$ is the true set of instances in the class c . The set H_c^T consists of instances that belong to class c . Therefore, $\{H_c^T\}_{c=1}^C$ should form a partition of all instances $\{1, \dots, K\}$. Let

$$P_{ic}^T = \Pr(i \in H_c^* | \mathcal{F}_T) = \Pr(\theta_{ic} \geq \theta_{i\tilde{c}}, \forall \tilde{c} \neq c | \mathcal{F}_T). \quad (54)$$

To maximize the right hand side of (53), we have

$$H_c^T = \{i : P_{ic}^T \geq P_{i\tilde{c}}^T, \forall \tilde{c} \neq c\}. \quad (55)$$

If there is i belongs to more than one H_c^T , we only assign it to the one with the smallest index c . The maximum conditional expected accuracy takes the form: $\sum_{i=1}^K (\max_{c \in \{1, \dots, C\}} P_{ic}^T)$.

Then the value function can be defined as:

$$V(S^0) \doteq \sup_{\pi} \mathbb{E}^{\pi} \left(\mathbb{E} \left(\sum_{i=1}^K \sum_{c=1}^C I(i \in H_c^T) I(i \in H_c^*) \middle| \mathcal{F}_T \right) \right) = \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{i=1}^K h(\mathbf{P}_i^T) \right),$$

where $\mathbf{P}_i^T = (P_{i1}^T, \dots, P_{iC}^T)$ and $h(\mathbf{P}_i^T) \doteq \max_{c \in \{1, \dots, C\}} P_{ic}^T$. Following Proposition 4, let $P_{ic}^t = \Pr(i \in H_c^* | \mathcal{F}_t)$ and $\mathbf{P}_i^t = (P_{i1}^t, \dots, P_{iC}^t)$, we define incremental reward function at each stage:

$$R(S^t, i_t) = \mathbb{E} (h(\mathbf{P}_{i_t}^{t+1}) - h(\mathbf{P}_{i_t}^t) | S^t, i_t).$$

The value function can be re-written as:

$$V(S^0) = G_0(S^0) + \sup_{\pi} \mathbb{E}^{\pi} \left(\sum_{t=0}^{T-1} R(S^t, i_t) \middle| S^0 \right),$$

where $G_0(S^0) = \sum_{i=1}^K h(\mathbf{P}_i^0)$. Since the reward function only depends on $S_{i_t}^t = \boldsymbol{\alpha}_{i_t}^t \in \mathbb{R}_+^C$, we can define the reward function in a more explicit way by defining:

$$R(\boldsymbol{\alpha}) = \sum_{c=1}^C \frac{\alpha_c}{\sum_{\tilde{c}=1}^C \alpha_{\tilde{c}}} h(I(\boldsymbol{\alpha} + \boldsymbol{\delta}_c)) - h(I(\boldsymbol{\alpha})).$$

Here δ_c be a row vector of length C with one at the c -th entry and zeros at all other entries; and $I(\alpha) = (I_1(\alpha), \dots, I_C(\alpha))$ where

$$I_c(\alpha) = \Pr(\theta_c \geq \theta_{\tilde{c}}, \forall \tilde{c} \neq c | \theta \sim \text{Dir}(\alpha)). \quad (56)$$

Therefore, we have $R(S^t, i_t) = R(\alpha_{i_t}^t)$.

To evaluate the reward $R(\alpha)$, the major bottleneck is how to compute $I_c(\alpha)$ efficiently. Directly taking the C -dimensional integration on the region $\{\theta_c \geq \theta_{\tilde{c}}, \forall \tilde{c} \neq c\} \cap \Delta_C$ will be computationally very expensive, where Δ_C denotes the C -dimensional simplex. Therefore, we propose a method to convert the computation of $I_c(\alpha)$ into a one-dimensional integration. It is known that to generate $\theta \sim \text{Dir}(\alpha)$, it is equivalent to generate $\{X_c\}_{c=1}^C$ with $X_c \sim \text{Gamma}(\alpha_c, 1)$ and let $\theta_c \equiv \frac{X_c}{\sum_{c=1}^C X_c}$. Then $\theta = (\theta_1, \dots, \theta_C)$ will follow $\text{Dir}(\alpha)$. Therefore, we have:

$$I_c(\alpha) = \Pr(X_c \geq X_{\tilde{c}}, \forall \tilde{c} \neq c | X_c \sim \text{Gamma}(\alpha_c, 1)). \quad (57)$$

It is easy to see that

$$\begin{aligned} I_c(\alpha) &= \int_{0 \leq x_1 \leq x_c} \cdots \int_{x_c \geq 0} \cdots \int_{0 \leq x_C \leq x_c} \prod_{c=1}^C f_{\text{Gamma}}(x_c; \alpha_c, 1) dx_1 \dots dx_C \\ &= \int_{x_c \geq 0} f_{\text{Gamma}}(x_c; \alpha_c, 1) \prod_{\tilde{c} \neq c} F_{\text{Gamma}}(x_c; \alpha_{\tilde{c}}, 1) dx_c, \end{aligned} \quad (58)$$

where $f_{\text{Gamma}}(x; \alpha_c, 1)$ is the density function of Gamma distribution with the parameter $(\alpha_c, 1)$ and $F_{\text{Gamma}}(x_c; \alpha_{\tilde{c}}, 1)$ is the CDF of Gamma distribution at x_c with the parameter $(\alpha_{\tilde{c}}, 1)$. In many softwares, $F_{\text{Gamma}}(x_c; \alpha_{\tilde{c}}, 1)$ can be calculated very efficiently without an explicit integration. Therefore, we can evaluate $I_c(\alpha)$ by performing only a one-dimensional numerical integration as in (58). We could also use Monte-Carlo approximation to further accelerate the computation in (58).

References

- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47:235–256, 2002.
- Y. Bachrach, T. Minka, J. Guiver, and T. Graepel. How to grade a test without knowing the answers - a Bayesian graphical model for adaptive crowdsourcing and aptitude testing. In *ICML*, 2012.
- M. J. Beal. *Variational Algorithms for Approximate Bayesian Inference*. PhD thesis, Gatsby Computational Neuroscience Unit, University College London, 2003.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2007.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.

- A. P. Dawid and A. M. Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society Series C*, 28:20–28, 1979.
- S. Ertekin, H. Hirsh, and C. Rudin. Wisely using a budget for crowdsourcing. Technical report, MIT, 2012.
- Eyal Even-Dar and Yishay Mansour. Convergence of optimistic and incremental Q-learning. In *NIPS*, 2001.
- P. Frazier, W. B. Powell, and S. Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM J. Control Optim.*, 47(5):2410–2439, 2008.
- C. Gao and D. Zhou. Minimax optimal convergence rates for estimating ground truth from crowdsourced labels. arXiv:1310.5764, 2013.
- Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall, 3rd edition, 2013.
- J. C. Gittins. *Multi-armed Bandit Allocation Indices*. John Wiley & Sons, 1989.
- S. S. Gupta and K. J. Miescke. Bayesian look ahead one stage sampling allocations for selection the largest normal mean. *J. of Stat. Planning and Inference*, 54(2):229–244, 1996.
- C. Ho, S. Jabbari, and J. W. Vaughan. Adaptive task assignment for crowdsourced classification. In *ICML*, 2013.
- P. G. Ipeirotis, F. Provost, V. S. Sheng, and J. Wang. Repeated labeling using multiple noisy label. *Data Mining and Knowledge Discovery*, 2013.
- T. Jaakkola and M. I. Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10:25–37, 2000.
- E. Kamar, S. Hacker, and E. Horvitz. Combing human and machine intelligence in large-scale crowdsourcing. In *AAMAS*, 2012.
- D. Karger, S. Oh, and D. Shah. Efficient crowdsourcing for multi-class labeling. In *ACM Sigmetrics*, 2013a.
- D. R. Karger, S. Oh, and D. Shah. Budget-optimal task allocation for reliable crowdsourcing systems. *Operations Research*, 62(1):1–24, 2013b.
- E. Kaufmann, O. Cappe, and A. Garivier. On Bayesian upper confidence bounds for bandit problems. In *AISTATS*, 2012.
- L. Li, W. Chu, J. Langford, and R. E. Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of International World Wide Web Conference*, 2010.
- C. Liu and Y. M. Wang. Truelabel + confusions: A spectrum of probabilistic models in analyzing multiple ratings. In *ICML*, 2012.

- Q. Liu, J. Peng, and A. Ihler. Variational inference for crowdsourcing. In *NIPS*, 2012.
- Andrew Y. Ng, Daishi Harada, and Stuart Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In *ICML*, 1999.
- J. Nino-Mora. Computing a classic index for finite-horizon bandits. *INFORMS Journal on Computing*, 23(2):254–267, 2011.
- R. D. Nowak. Noisy generalized binary search. In *NIPS*, 2009.
- J. Paisley, D. Blei, and M. Jordan. Variational bayesian inference with stochastic search. In *ICML*, 2012.
- W. B. Powell. *Approximate Dynamic Programming: solving the curses of dimensionality*. John Wiley & Sons, 2007.
- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 2005.
- V. C. Raykar, S. Yu, L. H. Zhao, G. H. Valadez, C. Florin, L. Bogoni, and L. Moy. Learning from crowds. *Journal of Machine Learning Research*, 11:1297–1322, 2010.
- C. P. Robert. *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2007.
- R. T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *J. of Banking and Finance*, 26:1443–1471, 2002.
- B. Settles. Active learning literature survey. Technical report, University of Wisconsin–Madison, 2009.
- R. Snow, B. O. Connor, D. Jurafsky, and A. Y. Ng. Cheap and fast - but is it good? evaluating non-expert annotations for natural language tasks. In *EMNLP*, 2008.
- István Szita and András Lőrincz. The many faces of optimism: a unifying approach. In *ICML*, 2008.
- J. Wang, P. G. Ipeirotis, and F. Provost. Quality-based pricing for crowdsourced workers. Technical report, New York University, 2013.
- P. Welinder, S. Branson, S. Belongie, and P. Perona. The multidimensional wisdom of crowds. In *NIPS*, 2010.
- J. Whitehill, P. Ruvolo, T. Wu, J. Bergsma, and J. R. Movellan. Whose vote should count more: Optimal integration of labels from labelers of unknown expertise. In *NIPS*, 2009.
- J. Xie and P. I. Frazier. Sequential bayes-optimal policies for multiple comparisons with a known standard. *Operations Research*, 61(5):1174–1189, 2013.
- Y. Yan, R. Rosales, G. Fung, and J. Dy. Active learning from crowds. In *ICML*, 2011.
- D. Zhou, S. Basu, Y. Mao, and J. Platt. Learning from the wisdom of crowds by minimax conditional entropy. In *NIPS*, 2012.