



Optimal clustering from noisy binary feedback

Kaito Ariu^{1,2} · Jungseul Ok³ · Alexandre Proutiere² · Seyoung Yun⁴

Received: 17 September 2021 / Revised: 5 February 2024 / Accepted: 15 February 2024 /
Published online: 22 March 2024
© The Author(s) 2024

Abstract

We study the problem of clustering a set of items from binary user feedback. Such a problem arises in crowdsourcing platforms solving large-scale labeling tasks with minimal effort put on the users. For example, in some of the recent reCAPTCHA systems, users clicks (binary answers) can be used to efficiently label images. In our inference problem, items are grouped into initially unknown non-overlapping clusters. To recover these clusters, the learner sequentially presents to users a finite list of items together with a question with a binary answer selected from a fixed finite set. For each of these items, the user provides a noisy answer whose expectation is determined by the item cluster and the question and by an item-specific parameter characterizing the *hardness* of classifying the item. The objective is to devise an algorithm with a minimal cluster recovery error rate. We derive problem-specific information-theoretical lower bounds on the error rate satisfied by any algorithm, for both uniform and adaptive (list, question) selection strategies. For uniform selection, we present a simple algorithm built upon the K-means algorithm and whose performance almost matches the fundamental limits. For adaptive selection, we develop an adaptive algorithm that is inspired by the derivation of the information-theoretical error lower bounds, and in turn allocates the budget in an efficient way. The algorithm learns to select items hard to cluster and relevant questions more often. We compare the performance of our algorithms with or without the adaptive selection strategy numerically and illustrate the gain achieved by being adaptive.

Keywords Online algorithm · Clustering · Community detection · Stochastic block model · Crowdsourcing

1 Introduction

Modern Machine Learning (ML) models require a massive amount of labeled data to be efficiently trained. Humans have been so far the main source of labeled data. This data collection is often tedious and very costly. Fortunately, most of the data can be simply labeled by non-experts. This observation is at the core of many crowdsourcing platforms such as reCAPTCHA, where users receive low or no payment. In these platforms, complex

Editor: Hendrik Blockeel.

Extended author information available on the last page of the article

labeling problems are decomposed into simpler tasks, typically questions with binary answers. In reCAPTCHAs, for example, the user is asked to click on images (presented in batches) that contain a particular object (a car, a road sign), and the system leverages users' answers to label images. As another motivating example, consider the task of classifying bird images. Users may be asked to answer binary questions like: "Is the bird grey?", "Does it have a circular tail fin?", "Does it have a pattern on its cheeks?", etc. Correct answers to those questions, if well-processed, may lead to an accurate bird classification and image labels. In both aforementioned examples, some images may be harder to label than others, e.g., due to the photographic environment, the birds' posture, etc. Some questions may be harder to answer than others, leading to a higher error rate. To build a reliable system, tasks/questions have to be carefully designed and selected, and user responses need to be smartly processed. Efficient systems must also learn the difficulty of the different tasks, and guess how informative they are when solving the complex labeling problem.

This paper investigates the design of such systems, tackling clustering problems that have to be solved using answers to binary questions. We incorporate a model that takes into consideration the varying difficulty levels or *heterogeneity* of clustering each item. We propose a full analysis of the problem, including information-theoretical limits that hold for *any* algorithm and novel algorithms with provable performance guarantees. Before giving a precise statement of our results, we provide a precise description of the problem setting and the statistical model dictating the way users answer. This model is inspired by models, such as the Dawid–Skene model (Dawid & Skene, 1979) successfully used in the crowdsourcing literature, see e.g., Khetan and Oh (2016) and references therein. However, to the best of our knowledge, this paper is the first to model and analyze clustering problems with binary feedback and accounting for item heterogeneity.

1.1 Problem setting and feedback model

Consider a large set \mathcal{I} of n items (e.g. images) partitioned into K disjoint unknown clusters $\mathcal{I}_1, \dots, \mathcal{I}_K$. Denote by $\sigma(i)$ the cluster of item i . To recover these hidden clusters, the learner gathers binary user feedback sequentially. Upon arrival, a user is presented a list of $w \geq 1$ items together with a question with a binary answer. The question is selected from a predefined finite set of cardinality L . The process of selecting the (list, question) pair for a given user can be carried out in either a nonadaptive or adaptive manner (in the latter case, the pair would depend on user feedback previously collected). Importantly, our model captures item heterogeneity: the difficulty of clustering items varies across items. We wish to devise algorithms recovering clusters as accurately as possible using the noisy binary answers collected from T users.

We use the following statistical model parametrized by a matrix $\mathbf{p} := (p_{k\ell})_{k \in [K], \ell \in [L]}$ ¹ with entries in $[0, 1]$ and by a vector $\mathbf{h} := (h_i)_{i \in \mathcal{I}} \in [1/2, 1]^n$. These parameters are (initially) unknown. When the t -th user is asked a question $\ell_t = \ell \in [L]$ for a set \mathcal{W}_t of $w \geq 1$ items, she provides noisy answers: for the item $i \in \mathcal{I}_k$ in the list \mathcal{W}_t , her answer $X_{i\ell_t}$ is $+1$ with probability $q_{i\ell} := h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}$, and -1 with probability $\bar{q}_{i\ell}$, where for brevity, \bar{x} denotes $1 - x$ for any $x \in [0, 1]$. Answers are independent across items and users. Our model is simple, but general enough to include as specific cases, crowdsourcing models recently investigated in the literature. For example, the model in Khetan and Oh (2016)

¹ Define for any integer $A \geq 1$, the set $[A] := \{1, \dots, A\}$.



Fig. 1 Images for Example 1

corresponds to our model with only one question ($L = 1$), two clusters ($K = 2$), and a question asked for a single item at a time ($w = 1$). Note that in our model, answers are collected from a very large set of users, and a given user is very unlikely to interact with the system several times. This justifies the fact that answers provided by the various users are statistically identical.

Item hardness. An important aspect of our model stems from the item-specific parameter h_i . It can be interpreted as the *hardness* of clustering item i , whereas $p_{k\ell}$ corresponds to a latent parameter related to question ℓ when asked for an item in cluster k . Note that when $h_i = 1/2$, $q_{i\ell} = 1/2$ irrespective of the cluster of item i . Hence any question ℓ on item i receives completely random responses, and this item cannot be clustered. Further note that intuitively, the larger the hardness parameter h_i of item i is, the easier it can be clustered. Indeed, when asking question ℓ , we can easily distinguish whether item i belongs to cluster k or k' if the difference between the corresponding parameters of user statistical answers $h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}$ and $h_i p_{k'\ell} + \bar{h}_i \bar{p}_{k'\ell}$ is large. This difference is $|p_{k\ell} - p_{k'\ell}|(2h_i - 1)$, an increasing function of h_i . We believe that introducing item heterogeneity is critical to obtain a realistic model (without h , all items from the same cluster would be exchangeable), but complicates the analysis. Most theoretical results on clustering or community detection do not account for this heterogeneity—refer to Sect. 2 for detail.

Illustrative Example. We introduce an example to illustrate the structure and characteristics of our model.

Example 1 Consider the task of classifying images into two types of birds: Mallards and Canadian Geese. Mallards (see Fig. 1a for an image), a type of duck, and Canadian Geese (see Fig. 1b for an image), which are not classified as ducks, present a unique classification challenge. In this case, $L = 1$, and the question posed to the users is: “Is the bird in the image a duck?”. We assign cluster 1 for the Mallard images and cluster 2 for the Canadian Goose images. Assume that $p_{11} = 0.8$ and $p_{21} = 0.3$: they are latent probabilities of answering yes to the question given an image of a Mallard and a Canadian Goose, respectively.

p_{11} and p_{21} represent the latent probabilities of answering yes to a question 1. These parameters also consider the scenario where a user, randomly selected from a large set, may not answer a question correctly due to a lack of knowledge or other reasons. For each image i , h_i indicates the difficulty of classification. For instance, when

the image is of Mallards (a type of duck), and the image is clear, the classification is relatively easy, and h_i is set to $h_i = 1$. Consequently, the probability of correctly identifying the Mallards is $q_{i1} = h_i p_{11} + \bar{h}_i \bar{p}_{11} = 1 \cdot 0.8 + 0 \cdot 0.2 = 0.8$. However, when another image j of Mallards is blurred due to poor lighting or other factors, and the classification difficulty is $h_j = 0.8$, the probability of correct classification decreases to $q_{j1} = h_j p_{11} + \bar{h}_j \bar{p}_{11} = 0.8 \cdot 0.8 + 0.2 \cdot 0.2 = 0.68$. As a result, the feedback obtained for image j is more ambiguous compared to that for image i , due to the increased difficulty in classification.

Assumption We make the following mild assumptions on our statistical model $\mathcal{M} := (\mathbf{p}, \mathbf{h})$. Define for each $k \in [K]$, $\mathbf{r}_k := (r_{k\ell})_{\ell \in [L]}$ with $r_{k\ell} := 2p_{k\ell} - 1$. Throughout the paper, $\|\cdot\|$ denotes the ℓ_∞ -norm, i.e., $\|\mathbf{x}\| = \max_i |x_i|$.

$$(A1) \quad h_* := \min_{i \in \mathcal{I}} (2h_i - 1) \in (0, 1), \quad (A2) \quad \exists \eta > 0, \eta \leq p_{k\ell} \leq 1 - \eta.$$

$$\rho_* := \min_{k \neq k'} \min_{c \geq 0} \|c\mathbf{r}_{k'} - \mathbf{r}_k\| > 0.$$

Assumption (A1) excludes the cases where clustering is impossible even if all parameters were accurately estimated. Indeed, when $h_* = 0$, there exists at least one item i which receives completely random responses for any question, i.e., $q_{i\ell} = 1/2$ for any $\ell \in [L]$. Observe that when $\rho_* = 0$, there exists $k \neq k'$ and $c \geq 0$ such that $2p_{k\ell} - 1 = c(2p_{k'\ell} - 1)$ for all $\ell \in [L]$. Then, for item $i \in \mathcal{I}_k$, we find $h' \in [1/2, 1]$ such that $2q_{i\ell} - 1 = (2h' - 1)(2p_{k'\ell} - 1)$. Items in the different clusters k and k' can have the same value of $q_{i\ell}$. As a consequence, from the answers, we cannot determine whether i is in cluster k or k' . In Example 1, $r_{11} = 0.6, r_{21} = -0.4$, and the value of ρ_* is $\rho_* = |0 \cdot 0.6 + 0.4| = 0.4$. Assumption (A2) states some homogeneity among the parameters of the clusters. It implies that $q_{i\ell} \in [\eta, 1 - \eta]$ for all $i \in \mathcal{I}$ and $\ell \in [L]$. Let Ω be the set of all models satisfying (A1) and (A2).

For convenience, we provide a table summarizing all the notations in Appendix A.

1.2 Main contributions

We study both nonadaptive and adaptive sequential (list, question) selection strategies. In the case of nonadaptive strategy, we assume that the selection of (list, question) pairs is *uniform* in the sense that the number of times a given question is asked for a given item is (roughly) $\lceil Tw/(nL) \rceil$. The objective is to devise a clustering algorithm taking as input the data collected over T users and returning estimated clusters as accurate as possible. When using adaptive strategies, the objective is to devise an algorithm that sequentially selects the (list, question) pairs presented to users, and that, after having collected answers from T users, returns accurate estimated clusters.

Our contributions are as follows. We first derive information-theoretical performance limits satisfied by any algorithm under uniform or adaptive sequential (list, question) selection strategy. We then propose a clustering algorithm that matches our limits order-wise in the case of uniform (list, question) selection. We further present a joint adaptive (list, question) selection strategy and clustering algorithm, and illustrate, using numerical experiments on both synthetic and real data, the advantage of being adaptive.

Fundamental limits. We provide a precise statement of our lower bounds on the cluster recovery error rate. These bounds are problem specific, i.e., they depend explicitly on the model $\mathcal{M} = (\mathbf{p}, \mathbf{h})$, and they will guide us in the design of algorithms.

(*Uniform selection*) In this case, we derive a clustering error lower bound for each individual item. Let π denote a clustering algorithm, and define the clustering error rate of item $i \in \mathcal{I}$ as $\epsilon_i^\pi(n, T) := \mathbb{P}[i \in \mathcal{E}^\pi]$, where \mathcal{E}^π denotes the set of mis-classified items under π . The latter set is defined as $\mathcal{E}^\pi := \cup_{k \in [K]} \mathcal{I}_k \setminus \mathcal{S}_{\gamma(k)}^\pi$, where $(\mathcal{S}_1^\pi, \dots, \mathcal{S}_K^\pi)$ denotes the output of π and γ is a permutation of $[K]$ minimizing the cardinality of $\cup_{k \in [K]} \mathcal{I}_k \setminus \mathcal{S}_{\zeta(k)}^\pi$ over all possible permutations ζ of $[K]$. When deriving problem-specific error lower bounds, we restrict our attention to so-called *uniformly good* algorithms. An algorithm π is *uniformly good* if for all $\mathcal{M} \in \Omega$ and $i \in \mathcal{I}$, $\epsilon_i^\pi(n, T) = o(1)$ as $T \rightarrow \infty$ under $T = \omega(n)$. We establish that for any $\mathcal{M} \in \Omega$ satisfying (A1) and (A2), under any uniformly good clustering algorithm π , as T grows large under $T = \omega(n)$, for any item i , we have:

$$\epsilon_i^\pi(n, T) \geq \exp\left(-\frac{T_W}{n} \mathcal{D}_{\mathcal{M}}^U(i)(1 + o(1))\right), \tag{1}$$

$$\text{where } \mathcal{D}_{\mathcal{M}}^U(i) := \min_{k' \neq \sigma(i)} \min_{h' \in [(h_{k'}+1)/2, 1]} \frac{1}{L} \sum_{\ell} \text{KL}(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, q_{i\ell}) > 0. \tag{2}$$

In the above definition of the *divergence* $\mathcal{D}_{\mathcal{M}}^U(i)$, $\text{KL}(a, b)$ is the Kullback–Leibler divergence between two Bernoulli distributions of means a and b ($\text{KL}(a, b) := a \log \frac{a}{b} + \bar{a} \log \frac{\bar{a}}{\bar{b}}$). Note that uniformly good algorithms actually exist (see Algorithm 1 presented in Sect. 4).

(*Adaptive selection*) We also provide clustering error lower bounds in the case the algorithm is also sequentially selecting (list, question) pairs in an adaptive manner. Note that here a lower bound cannot be derived for each item individually, say item i , since an adaptive algorithm could well select this given item often so as to get no error when returning its cluster. Instead we provide a lower bound for the cluster recovery error rate averaged over all items, i.e., $\epsilon^\pi(n, T) := \frac{1}{n} \sum_{i \in \mathcal{I}} \epsilon_i^\pi(n, T)$. Under any uniformly good joint (list, question) selection and clustering algorithm π , as T grows large under $T = \omega(n)$, we have:

$$\epsilon^\pi(n, T) \geq \exp\left(-\frac{T_W}{n} \tilde{\mathcal{D}}_{\mathcal{M}}^A(1 + o(1))\right), \tag{3}$$

$$\text{where } \tilde{\mathcal{D}}_{\mathcal{M}}^A := \max_{\mathbf{y} \in \mathcal{Y}(n)} -\frac{n}{T_W} \log\left(\frac{1}{n} \sum_{i=1}^n \exp\left(-\frac{T_W}{n} \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})\right)\right), \tag{4}$$

$$\mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y}) := \min_{j: \sigma(j) \neq \sigma(i)} \sum_{\ell} (y_{j\ell} \text{KL}(q_{j\ell}, q_{i\ell}) + y_{i\ell} \text{KL}(q_{i\ell}, q_{j\ell})), \tag{5}$$

$$\text{and } \mathcal{Y}(n) := \left\{ \mathbf{y} \in [0, 1]^{n \times L} : \sum_{i \in \mathcal{I}, \ell \in [L]} y_{i\ell} = n \right\}. \tag{6}$$

In the above lower bound, the vector \mathbf{y} encodes the expected numbers of times the various questions are asked for each item. Specifically, as shown later, $y_{i\ell} \frac{T_W}{n}$ can be interpreted expected number of times the question ℓ is asked for the item i . Maximizing over \mathbf{y} in (4)

hence corresponds to an optimal (list, question) selection strategy, and to the minimal error rate. Further interpretations and discussions of the divergences $\mathcal{D}_{\mathcal{M}}^U(i)$ and $\mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})$ are provided later in the paper.

Algorithms. We develop algorithms with both uniform and adaptive (list, question) selection strategies.

(Uniform selection) In this case, for each item i and based on the collected answers, we build a normalized vector (of dimension L) that concentrates (when T is large) around a vector depending on the cluster id $\sigma(i)$ only. Our algorithm applies a K-means algorithm to these vectors (with an appropriate initialization) to reconstruct the clusters. We are able to establish that the algorithm almost matches our fundamental limits. More precisely, when $T = \omega(n)$ and $T = o(n^2)$, under our algorithm, we have, for some absolute constant $C > 0$,

$$\varepsilon_i^\pi(n, T) \leq \exp\left(-C(2h_i - 1)^2 \rho_*^2 \frac{Tw}{Ln}(1 + o(1))\right). \quad (7)$$

The above error rate has an optimal scaling in T , w , L , n . By deriving upper and lower bounds on $\mathcal{D}_{\mathcal{M}}^U(i)$, we further show that the scaling is also optimal in $(2h_i - 1)^2$ and almost in ρ_* (see Assumption (A1)).

(Adaptive selection) The design of our adaptive algorithm is inspired by the information-theoretical lower bounds. The algorithm periodically updates estimates of the model parameters, and of the clusters. Based on these estimates, we further estimate lower bounds on the probabilities to misclassify every item. The items we select are those with the highest lower bounds (the items that are most likely to be misclassified); we further select the question that would be the most informative about these items. We believe that our algorithm should approach the minimal possible error rate (since it follows the optimal (list, question) selection strategy). Our numerical experiments suggest that the adaptive algorithm significantly outperforms algorithms with uniform (list, question) selection strategy, especially when items have very heterogenous hardnesses.

2 Related work

To our knowledge, the model proposed in this paper has been neither introduced nor analyzed in previous work. The problem has similarities with crowdsourced classification problems with a very rich literature (Dawid & Skene, 1979; Raykar et al., 2010; Karger et al., 2011; Zhou et al., 2012; Ho et al., 2013; Long et al., 2013; Zhang et al., 2014; Gao et al., 2016; Ok et al., 2016) (Dawid–Skene model and its various extensions without clustered structure), Vinayak and Hassibi (2016) and Gomes et al. (2011) (Clustering without item heterogeneity). However, our model has clear differences. For instance, if we draw a parallel between our model and that considered in Khetan and Oh (2016), there tasks correspond to our items, and there are only two clusters of tasks. More importantly, the statistics of the answers for a particular task do not depend on the true cluster of the task since the ground truth is defined by the majority of answers given by the various users. Our results also differ from those in the crowdsourcing literature from a methodological perspective. In this literature, fundamental limits are rarely investigated, and if they are, they are in the minimax sense by postulating the worst parameter setting (e.g., Zhang et al. 2014; Khetan and Oh 2016; Gao et al. 2016) or it is problem-specific but without quantifying of the error rate (e.g., Ok et al. 2016). Here we derive more precise problem-specific lower bounds on the error rate, i.e., we provide minimum clustering error

rates given the model parameters (p, h) . Further note that most of the classification tasks studied in the literature are simple (can be solved using a single binary question).

Our problem also resembles cluster inference problems in the celebrated Stochastic Block Model (SBM), see (Abbe, 2018) for a recent survey. Plain SBM models, however, assume that the statistics of observations for items in the same cluster are identical (there are no items harder to cluster than others, this corresponds to $h_i = 1, \forall i \in \mathcal{I}$ in our model), and observations are typically not operated in an adaptive manner. The closest work in the context of SBM to ours is the analysis of the so-called Degree-Corrected SBM, where each node is associated with an average degree quantifying the number of observations obtained for this node. The average degree then replaces our hardness parameter h_i for item i . In Gao et al. (2018), the authors study the Degree-Corrected SBM, but deal with minimax performance guarantees only, and non-adaptive sampling strategies.

3 Information-theoretical limits

3.1 Uniform selection strategy

Recall that an algorithm π is *uniformly good* if for all $\mathcal{M} \in \Omega$ and $i \in \mathcal{I}$, $\epsilon_i^\pi(n, T) = o(1)$ as $T \rightarrow \infty$ under $T = \omega(n)$. Assumptions (A1) and (A2) ensure the existence of uniformly good algorithms. The algorithm we present in Sect. 4 is uniformly good under these assumptions. The following theorem provides a lower bound on the error rate of uniformly good algorithms.

Theorem 1 *If an algorithm π with uniform selection strategy is uniformly good, then for any $\mathcal{M} \in \Omega$ satisfying (A1) and (A2), under $T = \omega(n)$, the following holds:*

$$\liminf_{T \rightarrow \infty} \frac{\frac{T_W}{n} \mathcal{D}_{\mathcal{M}}^U(i)}{\log(1/\epsilon_i^\pi(n, T))} \geq 1, \quad \forall i \in \mathcal{I}.$$

The proof of Theorem 1 will be presented later in this section. Theorem 1 implies that the global error rate of any uniformly good algorithm satisfies:

$$\epsilon^\pi(n, T) \geq \exp\left(-\frac{T_W}{n} \tilde{\mathcal{D}}_{\mathcal{M}}^U(1 + o(1))\right),$$

$$\text{where } \tilde{\mathcal{D}}_{\mathcal{M}}^U := -\frac{n}{T_W} \log\left(\frac{1}{n} \sum_{i \in \mathcal{I}} \exp\left(-\frac{T_W}{n} \mathcal{D}_{\mathcal{M}}^U(i)\right)\right)$$

$$\text{and } \mathcal{D}_{\mathcal{M}}^U(i) := \min_{k' \neq \sigma(i)} \min_{h' \in ((h_i + 1)/2, 1]} \frac{1}{L} \sum_{\ell} \text{KL}(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, q_{i\ell}) > 0.$$

Divergence $\mathcal{D}_{\mathcal{M}}^U(i)$ and its properties. The *divergence* $\mathcal{D}_{\mathcal{M}}^U(i)$, defined in Section 1, quantifies the hardness of classifying item i . This divergence naturally appears in the change-of-measure argument used to establish Theorem 1. To get a better understanding of $\mathcal{D}_{\mathcal{M}}^U(i)$, and in particular to assess its dependence on the various system parameters, we provide the following useful upper and lower bounds, proved in Appendix B:

Proposition 1 *Fix $i \in \mathcal{I}$. Let k' be such that:*

$$\mathcal{D}_{\mathcal{M}}^U(i) = \min_{h' \in [(h_*+1)/2, 1]} \frac{1}{L} \sum_{\ell} \text{KL}(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, q_{i\ell}).$$

Then, we have:

$$\begin{aligned} \mathcal{D}_{\mathcal{M}}^U(i) &\leq \frac{1}{2L\eta} (2h_i - 1)^2 \min_{\frac{h_*}{2h_i-1} \leq c \leq \frac{1}{2h_i-1}} \|c\mathbf{r}_{k'} - \mathbf{r}_{\sigma(i)}\|_2^2, \\ \mathcal{D}_{\mathcal{M}}^U(i) &\geq \frac{1}{2L} (2h_i - 1)^2 \min_{\frac{h_*}{2h_i-1} \leq c \leq \frac{1}{2h_i-1}} \|c\mathbf{r}_{k'} - \mathbf{r}_{\sigma(i)}\|_2^2. \end{aligned} \tag{8}$$

Note that $\mathcal{D}_{\mathcal{M}}^U(i)$ vanishes as h_i goes to $1/2$, which makes sense since for $h_i \approx 1/2$, item i is very hard to cluster. We also have $\mathcal{D}_{\mathcal{M}}^U(i) = 0$ when $\min_{\frac{h_*}{2h_i-1} \leq c \leq \frac{1}{2h_i-1}} \|c\mathbf{r}_{k'} - \mathbf{r}_{\sigma(i)}\|_2^2 = 0$. In this case, $\rho_* = 0$ and there exists $h' \in [(1 + h_*)/2, 1]$ such that for some $k' \neq \sigma(i)$, $2q_{i\ell} - 1 = (2h' - 1)(2p_{k'\ell} - 1)$ for all $\ell \in [L]$, so that clustering item i is impossible.

Application to the simpler model of Khetan and Oh (2016). Consider a model with a single question and two clusters of items. From Theorem 1, we can recover an asymptotic version of Theorem 2.4. in Khetan and Oh (2016).

Corollary 1 *Let $L = 1, K = 2, \mathbf{p} = (p_{11}, p_{21})$, and $w = 1$. If an algorithm π with uniform selection strategy is uniformly good, whenever \mathcal{M} satisfies (A1) and (A2), under $T = \omega(n)$, we have:*

$$\liminf_{T \rightarrow \infty} \frac{T C (2h_i - 1)^2 (p_{11} - p_{21})^2}{\log(1/\varepsilon_i^\pi(n, T))} \geq 1, \quad \forall i \in \mathcal{I},$$

where $C > 0$ is an absolute constant.

The proof of Corollary 1 is presented in Appendix D. Corollary 1 implies,

$$\varepsilon_i^\pi(n, T) \geq \exp\left(-\frac{T}{n} C (2h_i - 1)^2 (p_{11} - p_{21})^2 (1 + o(1))\right)$$

as $T \rightarrow \infty$ under $T = \omega(n)$. Smaller h_i and $|p_{11} - p_{21}|$ imply item i is harder to classify. Note that Theorem 2.4. in Khetan and Oh (2016) (corresponds to $p_{21} = 1 - p_{11}$ in our Corollary 1) provides a minimax lower bound whereas our result is problem-specific and hence more precise. Note that Corollary 1 also applies directly to Example 1 mentioned in Introduction. The lower bound on the error probability for each item i scales as $\exp(-c \frac{T}{n} (2h_i - 1)^2)$ with some constant $c > 0$.

Proof of Theorem 1 The proof leverages change-of-measure arguments, as those used in the classical multi-armed bandit problem (Lai & Herbert, 1985) or the Stochastic Block Model (Yun & Proutiere, 2016). Here the proof is however complicated by the fact that we wish a lower bound on the error rate for clustering each item.

Let π denote a uniformly good algorithm with uniform selection strategy, and let $\mathcal{M} \in \Omega$ be a model satisfying Assumptions (A1) and (A2). In our change-of-measure, we denote by \mathcal{M} the original model and by \mathcal{N} a perturbed model. Fix $i \in \mathcal{I}$, where $\sigma(i) = k$. Let $k' \in [K], h' \in [(h_* + 1)/2, 1]$ denote the minimizers for the optimization problem leading to $\mathcal{D}_{\mathcal{M}}^U(i)$, i.e.,

$$\mathcal{D}_{\mathcal{M}}^U(i) = \frac{1}{L} \sum_{\ell=1}^L \text{KL}(h'p_{k'\ell} + \bar{h}'\bar{p}_{k'\ell}, q_{i\ell}) > 0.$$

For these choices of i, k' , and h' , we construct the perturbed model \mathcal{N} as follows. Under \mathcal{N} , all responses for items different than i are generated as under \mathcal{M} . The responses for i under \mathcal{N} are generated as if i was in cluster k' and had difficulty h' . We can write the log-likelihood ratio of the observation under \mathcal{N} to that under \mathcal{M} as follows:

$$\mathcal{L} = \sum_{i=1}^T \mathbb{1}[i \in \mathcal{W}_i] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell t} = +1] \log \frac{q'_{\ell}}{q_{i\ell}} + \mathbb{1}[X_{i\ell t} = -1] \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right). \tag{9}$$

where we let $\mathbf{q}' := (q'_{\ell})_{\ell \in [L]}$ with $q'_{\ell} = h'p_{k'\ell} + \bar{h}'\bar{p}_{k'\ell}$.

Let $\mathbb{P}_{\mathcal{N}}$ and $\mathbb{E}_{\mathcal{N}}$ (resp. $\mathbb{P}_{\mathcal{M}} = \mathbb{P}$ and $\mathbb{E}_{\mathcal{M}} = \mathbb{E}$) denote, respectively, the probability measure and the expectation under \mathcal{N} (resp. \mathcal{M}). Using the construction of \mathcal{N} , a change-of-measure argument provides us with a connection between the error rate on item i under \mathcal{M} and the mean and the variance of \mathcal{L} under \mathcal{N} :

$$\log(1/\varepsilon_i^{\pi}(n, T)) - \log 4 \leq \mathbb{E}_{\mathcal{N}}[\mathcal{L}] + \sqrt{2\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]}. \tag{10}$$

□

Proof of (10) The distribution of the log-likelihood \mathcal{L} under \mathcal{N} satisfies: for any $g \geq 0$,

$$\begin{aligned} \mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq g\} &= \mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\} + \mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq g, i \notin \mathcal{E}^{\pi}\} \\ &\leq \mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\} + \mathbb{P}_{\mathcal{N}}\{i \notin \mathcal{E}^{\pi}\}. \end{aligned} \tag{11}$$

Using the definition (9) of the log-likelihood ratio, we bound the first term in (11) as follows:

$$\begin{aligned} \mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\} &= \int_{\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\}} d\mathbb{P}_{\mathcal{N}} \\ &= \int_{\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\}} \exp(\mathcal{L}) d\mathbb{P}_{\mathcal{M}} \\ &\leq \exp(g) \mathbb{P}_{\mathcal{M}}\{\mathcal{L} \leq g, i \in \mathcal{E}^{\pi}\} \\ &\leq \exp(g) \varepsilon_i^{\pi}(n, T). \end{aligned} \tag{12}$$

To bound the second term in (11), note that $(2h' - 1)$ is a strictly positive constant.² Hence, the perturbed model \mathcal{N} satisfies (A1). By the definition of the uniformly good algorithm, we have $\mathbb{P}_{\mathcal{N}}\{i \notin \mathcal{S}_{k'}^{\pi}\} = o(1)$. Hence:

$$\mathbb{P}_{\mathcal{N}}\{i \notin \mathcal{E}^{\pi}\} = \mathbb{P}_{\mathcal{N}}\{i \in \mathcal{S}_k^{\pi}\} \leq \mathbb{P}_{\mathcal{N}}\{i \notin \mathcal{S}_{k'}^{\pi}\} \leq \frac{1}{4}. \tag{13}$$

Combining (11), (12) and (13) with $g = -\log(4\varepsilon_i^{\pi}(n, T))$, we have

² This is true as h' is optimized over $[(h_* + 1)/2, 1]$.

$$\mathbb{P}_{\mathcal{N}}\{\mathcal{L} \leq -\log \varepsilon_i^\pi(n, T) - \log 4\} \leq \frac{1}{2}. \tag{14}$$

Using Chebyshev’s inequality, we obtain:

$$\mathbb{P}_{\mathcal{N}}\left\{\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}] \geq \sqrt{2\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]}\right\} \leq \frac{\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]}{\left(\sqrt{2\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]}\right)^2} = \frac{1}{2},$$

which implies $\mathbb{P}_{\mathcal{N}}\left\{\mathcal{L} \leq \mathbb{E}_{\mathcal{N}}[\mathcal{L}] + \sqrt{2\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]}\right\} \geq \frac{1}{2}$. Combining this result to (14) implies the claim (10). □

Next, Lemma 1 provides the upper bound on mean and variance of \mathcal{L} under the model \mathcal{N} .

Lemma 1 *Assume that (A2) holds. For i, i' such that $\sigma(i) = k \neq k' = \sigma(i')$, under the uniform selection strategy, we have*

$$\begin{aligned} \mathbb{E}_{\mathcal{N}}[\mathcal{L}] &= \frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^U(i), \quad \text{and} \\ \mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2] &\leq \frac{Tw}{Ln} \log\left(\frac{1}{\eta}\right) \sum_{\ell=1}^L \text{KL}(q'_{\ell}, q_{i\ell}) + \sqrt{\text{KL}(q_{i'\ell}, q_{i\ell})}. \end{aligned}$$

The proof of this lemma is presented to Appendix C. Note that in view the above lemma, the r.h.s. of (10) is asymptotically dominated by $\mathbb{E}_{\mathcal{N}}[\mathcal{L}]$, since $\mathbb{E}_{\mathcal{N}}[\mathcal{L}] = \Omega(T/n)$ and $\sqrt{\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2]} = O(\sqrt{T/n})$. Thus, Theorem 1 follows from the claim in (10) and Lemma 1. □

3.2 Adaptive selection strategy

The derivation of a lower bound for the error rate under adaptive (list, pair) selection strategies is similar:

Theorem 2 *For any $\mathcal{M} \in \Omega$ satisfying (A1) and (A2), and for any uniformly good algorithm π with possibly adaptive (list, question) selection strategy, under $T = \omega(n)$, we have:*

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{\frac{Tw}{n} \tilde{\mathcal{D}}_{\mathcal{M}}^A}{\log(1/\varepsilon^\pi(n, T))} &\geq 1, \\ \text{where } \tilde{\mathcal{D}}_{\mathcal{M}}^A &:= \max_{\mathbf{y} \in \mathcal{Y}(n)} -\frac{n}{Tw} \log\left(\frac{1}{n} \sum_{i=1}^n \exp\left(-\frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})\right)\right), \\ \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y}) &:= \min_{j: \sigma(j) \neq \sigma(i)} \sum_{\ell} (y_{j\ell} \text{KL}(q_{j\ell}, q_{i\ell}) + y_{i\ell} \text{KL}(q_{i\ell}, q_{j\ell})), \\ \text{and } \mathcal{Y}(n) &:= \left\{ \mathbf{y} \in [0, 1]^{n \times L} : \sum_{i \in \mathcal{I}, \ell \in [L]} y_{i\ell} = n \right\}. \end{aligned}$$

Theorem 2 implies $\epsilon^\pi(n, T) \geq \exp(-\frac{Tw}{n} \widetilde{D}_{\mathcal{M}}^A(1 + o(1)))$.

Proof of Theorem 2 Again we use a change-of-measure argument, where we swap two items from different clusters. First, we prove the lower bound for the error rate of a fixed item i . Fix $i \in \mathcal{I}$, let j be an item satisfying $\sigma(j) \neq \sigma(i)$ and let

$$D_{\mathcal{M}}^A(i, \mathbf{y}) = \sum_{\ell=1}^L (y_{j\ell} \text{KL}(q_{j\ell}, q_{i\ell}) + y_{i\ell} \text{KL}(q_{i\ell}, q_{j\ell})).$$

$D_{\mathcal{M}}^A(i, \mathbf{y})$ is the value of the optimization problem:

$$\min_{j: \sigma(j) \neq \sigma(i)} \sum_{\ell} (y_{j\ell} \text{KL}(q_{j\ell}, q_{i\ell}) + y_{i\ell} \text{KL}(q_{i\ell}, q_{j\ell})).$$

Consider a perturbed model \mathcal{N} , in which items except i and j have the same response statistics as under \mathcal{M} , and in which item i behaves as item j , and item j behaves as item i . Let $\mathbb{P}_{\mathcal{N}}$ and $\mathbb{E}_{\mathcal{N}}$ denote, respectively, the probability measure and the expectation under \mathcal{N} . The log-likelihood ratio of the responses under \mathcal{N} and under \mathcal{M} is:

$$\begin{aligned} \mathcal{L} &= \sum_{t=1}^T \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell t} = +1] \log \frac{q_{j\ell}}{q_{i\ell}} + \mathbb{1}[X_{i\ell t} = -1] \log \frac{\bar{q}_{j\ell}}{\bar{q}_{i\ell}} \right) \\ &+ \sum_{t=1}^T \mathbb{1}[j \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{j\ell t} = +1] \log \frac{q_{i\ell}}{q_{j\ell}} + \mathbb{1}[X_{j\ell t} = -1] \log \frac{\bar{q}_{i\ell}}{\bar{q}_{j\ell}} \right). \end{aligned} \tag{15}$$

The mean and variance of \mathcal{L} under \mathcal{N} are:

$$\begin{aligned} \mathbb{E}_{\mathcal{N}}[\mathcal{L}] &= \sum_{\ell=1}^L (\mathbb{E}_{\mathcal{N}}[Y_{i\ell}] \text{KL}(q_{j\ell}, q_{i\ell}) + \mathbb{E}_{\mathcal{N}}[Y_{j\ell}] \text{KL}(q_{i\ell}, q_{j\ell})) \\ &= \frac{Tw}{n} \sum_{\ell=1}^L (y_{j\ell} \text{KL}(q_{j\ell}, q_{i\ell}) + y_{i\ell} \text{KL}(q_{i\ell}, q_{j\ell})) \\ &= \frac{Tw}{n} D_{\mathcal{M}}^A(i, \mathbf{y}), \\ \mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2] &\leq \log \left(\frac{1}{\eta} \right) \sum_{\ell=1}^L \mathbb{E}_{\mathcal{N}}[Y_{i\ell}] \left(\text{KL}(q_{j\ell}, q_{i\ell}) + \sqrt{\text{KL}(q_{j\ell}, q_{i\ell})} \right) \\ &\quad + \log \left(\frac{1}{\eta} \right) \sum_{\ell=1}^L \mathbb{E}_{\mathcal{N}}[Y_{j\ell}] \left(\text{KL}(q_{i\ell}, q_{j\ell}) + \sqrt{\text{KL}(q_{i\ell}, q_{j\ell})} \right) \\ &= \frac{Tw}{n} \log \left(\frac{1}{\eta} \right) \sum_{\ell=1}^L y_{j\ell} \left(\text{KL}(q_{j\ell}, q_{i\ell}) + \sqrt{\text{KL}(q_{j\ell}, q_{i\ell})} \right) \\ &\quad + \frac{Tw}{n} \log \left(\frac{1}{\eta} \right) \sum_{\ell=1}^L y_{i\ell} \left(\text{KL}(q_{i\ell}, q_{j\ell}) + \sqrt{\text{KL}(q_{i\ell}, q_{j\ell})} \right), \end{aligned}$$

using a slight modification of Lemma 1. By a similar argument as that used in the proof of Theorem 1, we get:

$$\log(1/\varepsilon_i^\pi(n, T)) - \log 4 \leq \mathbb{E}_{\mathcal{M}}[\mathcal{L}] + \sqrt{2\mathbb{E}_{\mathcal{M}}[(\mathcal{L} - \mathbb{E}_{\mathcal{M}}[\mathcal{L}])^2]}, \tag{16}$$

as is in (10). Note that the r.h.s. of (16) is asymptotically dominated by $\mathbb{E}_{\mathcal{M}}[\mathcal{L}]$ as $\mathbb{E}_{\mathcal{M}}[\mathcal{L}] = \Omega(T/n)$ and $\sqrt{\mathbb{E}_{\mathcal{M}}[(\mathcal{L} - \mathbb{E}_{\mathcal{M}}[\mathcal{L}])^2]} = O(\sqrt{T/n})$. That is,

$$\mathbb{E}_{\mathcal{M}}[\mathcal{L}] + \sqrt{2\mathbb{E}_{\mathcal{M}}[(\mathcal{L} - \mathbb{E}_{\mathcal{M}}[\mathcal{L}])^2]} = \frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})(1 + o(1)).$$

We deduce that $\varepsilon_i^\pi(n, T) \geq \exp\left(-\frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})(1 + o(1))\right)$. Thus, from the definition of $\tilde{D}_{\mathcal{M}}^A$, we have:

$$\begin{aligned} \varepsilon^\pi(n, T) &= \frac{1}{n} \sum_{i \in \mathcal{I}} \varepsilon_i^\pi(n, T) \\ &= \frac{1}{n} \sum_{i \in \mathcal{I}} \exp\left(-\frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})(1 + o(1))\right) \\ &\geq \exp\left(-\frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^A(1 + o(1))\right), \end{aligned}$$

Taking the logarithm of the previous inequality, we conclude the proof. □

4 Algorithms

In this section, we describe our algorithms for both uniform and adaptive (list, question) selection strategies.

Algorithm 1 Uniform (list, question) selection.

Input: K, T

User response collection:

Question ℓ is asked to each item i for $\tau = \lfloor \frac{Tw}{nL} \rfloor$ times.

For all $\ell \in [L]$ and $i \in \mathcal{I}$, $x_{i\ell} \leftarrow \sum_t^T \mathbb{1}[i \in \mathcal{W}_t, \ell_t = \ell, X_{i\ell t} = 1]$ and $\hat{q}_{i\ell} \leftarrow x_{i\ell}/\tau$.

Normalization:

$$\hat{\mathbf{r}}^i \leftarrow \frac{2\hat{\mathbf{q}}_i - 1}{\|2\hat{\mathbf{q}}_i - 1\|} \text{ for all } i \in \mathcal{I}$$

K-means clustering:

$$T_i \leftarrow \{j \in \mathcal{I} : \|\hat{\mathbf{r}}^j - \hat{\mathbf{r}}^i\|^2 \leq (\frac{n}{T})^{\frac{1}{2}}\} \text{ for each } i \in \mathcal{I}$$

$$S_0 \leftarrow \emptyset$$

for $k = 1$ **to** K **do**

$$i_k^* \leftarrow \arg \max_{i \in \mathcal{I}} |T_i \setminus \bigcup_{k'=0}^{k-1} S_{k'}|$$

$$S_k \leftarrow T_{i_k^*} \setminus \bigcup_{k'=0}^{k-1} S_{k'}$$

$$\xi_k \leftarrow (\sum_{i \in S_k} \hat{\mathbf{r}}^i) / |S_k|$$

end for

For each $i \in \mathcal{I} \setminus \bigcup_{k=1}^K S_k$ (or each $i \in \mathcal{I}$), place i in S_{k^*} , where $k^* = \arg \min_k \|\xi_k - \hat{\mathbf{r}}^i\|$

Output: $\{S_k\}_{k \in [K]}$

4.1 Uniform selection strategy

In this case, we assume that with a budget of T users, each item receives the same amount of answers for each question. After gathering these answers, we have to exploit the data to estimate the clusters. To this aim, we propose an extension of the K-means clustering algorithm, that efficiently leverages the problem structure. The pseudo-code of the algorithm is presented in Algorithm 1.

The algorithm first estimates the parameters $q_{\ell i}$: the estimator $\hat{q}_{\ell i}$ just counts the number of times item i has received a positive answer for question ℓ . We denote by $\hat{\mathbf{q}}_i = (\hat{q}_{\ell i})_{\ell}$ the resulting vector. By normalizing the vector $2\hat{\mathbf{q}}_i - \mathbf{1}$, we can decouple the nonlinear relationship between \mathbf{q} , \mathbf{h} and \mathbf{p} . Let $\hat{\mathbf{r}}^i = \frac{2\hat{\mathbf{q}}_i - \mathbf{1}}{\|2\hat{\mathbf{q}}_i - \mathbf{1}\|}$ be the normalized vector. Then, $\hat{\mathbf{r}}^i$ concentrates around $\tilde{\mathbf{r}}_{\sigma(i)} := \mathbf{r}_{\sigma(i)} / \|\mathbf{r}_{\sigma(i)}\|$. Importantly, the normalized vector $\tilde{\mathbf{r}}_{\sigma(i)}$ does not depend on h_i but on the cluster index $\sigma(i)$ only. The algorithm exploits this observation, and applies the K-means algorithm to cluster the vectors $\hat{\mathbf{r}}^i$. By analyzing how $\hat{\mathbf{r}}^i$ concentrates around $\tilde{\mathbf{r}}_{\sigma(i)}$ and by applying the results to our properly tuned algorithm (decision thresholds), we establish the following theorem.

Theorem 3 *Assume $T = \omega(n)$ and $T = o(n^2)$. Under Algorithm 1, we have,*

$$\varepsilon_i(n, T) \leq \exp\left(-\frac{(2h_i - 1)^2 \rho_*^2 Tw}{200 Ln} (1 + o(1))\right). \tag{17}$$

We will present the proof of Theorem 3 later in this section. In view of Proposition 1 and the lower bounds derived in the previous section, we observe that the exponent for the mis-classification error of item i has the correct dependence in Tw/Ln and the tightest possible scaling in the hardness of the item, namely $(2h_i - 1)^2$. Also note that using Proposition 1, the equivalence between the ℓ_{∞} -norm and the Euclidean norm, and (A1), we have: $\mathcal{D}_{\mathcal{M}}^U(i) \geq C \frac{(2h_i - 1)^2}{L} \rho_*^2$, for some absolute constant $C > 0$. Hence, Algorithm 1 has a performance scaling optimally w.r.t. all the model parameters.

The computational complexity of Algorithm 1 is $\mathcal{O}(n^2)$. By choosing a small ($\log n$) subset of items (and not all the items in \mathcal{I}) to compute centroids (T_i), it is possible to reduce the computational complexity to $\mathcal{O}(n \log n)$. This would not affect the performance of the algorithm in practice, but would result in worse performance guarantees.

Proof of Theorem 3 In this proof, we let $\tau = \lfloor \frac{Tw}{nL} \rfloor$ be the number of times question ℓ is asked for item i . We also denote by $\boldsymbol{\alpha} := (\alpha_1, \dots, \alpha_K)$ the fractions of items that are in the various clusters, i.e., $|\mathcal{I}_k| = \alpha_k n$. Without loss of generality, and to simplify the notation, we assume that the set of misclassified items is $\mathcal{E} = \cup_{k=1}^K (\mathcal{I}_k \setminus \mathcal{S}_k)$, where recall that $\{\mathcal{S}_k\}_{k \in [K]}$ is the output of the algorithm (i.e., the permutation γ in the definition of this set is just the identity).

The proof proceeds in three steps: (i) we first decompose the probability of clustering error for item i , using the design of the algorithm and Assumptions (A1) and (A2). We show that this probability can be upper bounded by the probabilities of events related to $\|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}_{\sigma(i)}\|$ and $\|\xi_k - \tilde{\mathbf{r}}_k\|$ for all k , where recall that $\tilde{\mathbf{r}}_k := \mathbf{r}_k / \|\mathbf{r}_k\|$. The remaining steps of the proof aim at bounding the probabilities of these events. Towards this objective, (ii) in the second step, we establish a concentration result on $\|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}_{\sigma(i)}\|$, and (iii) the last step upper bound $\|\xi_k - \tilde{\mathbf{r}}_k\|$.

Step 1. Error probability decomposition. The algorithm classifies item i to the cluster k minimizing the distance between \hat{r}^i and ξ_k . As a consequence, we have:

$$\begin{aligned} \mathbb{P}\{i \in \mathcal{E}\} &= \mathbb{P}\left\{\|\hat{r}^i - \xi_{\sigma(i)}\| \geq \min_{k' \neq \sigma(i)} \|\hat{r}^i - \xi_{k'}\|\right\} \\ &\leq \mathbb{P}\left\{\|\hat{r}^i - \tilde{r}_{\sigma(i)}\| + \|\tilde{r}_{\sigma(i)} - \xi_{\sigma(i)}\| \geq \min_{k' \neq \sigma(i)} \{\|\hat{r}^i - \tilde{r}_{k'}\| - \|\tilde{r}_{k'} - \xi_{k'}\|\}\right\} \\ &\leq \mathbb{P}\left\{\exists k' \neq \sigma(i) : \|\tilde{r}_{\sigma(i)} - \xi_{\sigma(i)}\| + \|\tilde{r}_{k'} - \xi_{k'}\| + 2\|\hat{r}^i - \tilde{r}_{\sigma(i)}\| \geq \|\tilde{r}_{\sigma(i)} - \tilde{r}_{k'}\|\right\}, \end{aligned}$$

where the two above inequalities are obtained by simply applying the triangle inequality. Now observe that in view of Assumptions (A1) and (A2), we have: for $k' \neq \sigma(i)$,

$$\|\tilde{r}_{\sigma(i)} - \tilde{r}_{k'}\| = \frac{1}{\|\mathbf{r}_{\sigma(i)}\|} \left\| \mathbf{r}_{\sigma(i)} - \frac{\|\mathbf{r}_{\sigma(i)}\|}{\|\mathbf{r}_{k'}\|} \mathbf{r}_{k'} \right\| \geq \frac{\rho_*}{\|\mathbf{r}_{\sigma(i)}\|} = \frac{\rho_*}{\|2\mathbf{p}_{\sigma(i)} - \mathbf{1}\|} \geq \frac{\rho_*}{|2\eta - 1|}.$$

We deduce that:

$$\begin{aligned} \mathbb{P}\{i \in \mathcal{E}\} &\leq \mathbb{P}\left\{\exists k' \neq \sigma(i) : \|\tilde{r}_{\sigma(i)} - \xi_{\sigma(i)}\| + \|\tilde{r}_{k'} - \xi_{k'}\| + 2\|\hat{r}^i - \tilde{r}_{\sigma(i)}\| \geq \frac{\rho_*}{\|2\mathbf{p}_{\sigma(i)} - \mathbf{1}\|}\right\}, \\ &\leq \mathbb{P}\left\{\|\hat{r}^i - \tilde{r}_{\sigma(i)}\| \geq \frac{\rho_*}{4\|2\mathbf{p}_{\sigma(i)} - \mathbf{1}\|} \text{ or } \max_{1 \leq k' \leq K} \|\xi_{k'} - \tilde{r}_{k'}\| \geq \frac{\rho_*}{4\|2\mathbf{p}_{\sigma(i)} - \mathbf{1}\|}\right\} \\ &\leq \underbrace{\mathbb{P}\left\{\|\hat{r}^i - \tilde{r}_{\sigma(i)}\| \geq \frac{\rho_*}{4\|2\mathbf{p}_{\sigma(i)} - \mathbf{1}\|}\right\}}_{\text{term (a)}} + \underbrace{\sum_{k'=1}^K \mathbb{P}\left\{\|\xi_{k'} - \tilde{r}_{k'}\| \geq \frac{\rho_*}{4|2\eta - 1|}\right\}}_{\text{term (b)}}. \end{aligned}$$

Step 2. Concentration of \hat{r}^i and upper bound on (a). We prove the following lemma, a consequence of a concentration result of \hat{q}_i : □

Lemma 2 Let $0 < \varepsilon \leq \frac{\|2\mathbf{q}_i - \mathbf{1}\|}{16}$, $\tilde{r}^i = \frac{\mathbf{r}_{\sigma(i)}}{\|\mathbf{r}_{\sigma(i)}\|}$, and $\tilde{r}_k = \frac{\mathbf{r}_k}{\|\mathbf{r}_k\|}$. For each $i \in \mathcal{I}$,

$$\|\hat{r}^i - \tilde{r}^i\| \leq \frac{5\varepsilon}{\|2\mathbf{q}_i - \mathbf{1}\|},$$

with probability at least $1 - 2L \exp(-2\tau\varepsilon^2)$.

The proof of Lemma 2 is presented at Appendix E. Note that by definition of ρ_* , we have:

$$(2h_i - 1)\rho_* \leq (2h_i - 1)\|\mathbf{r}_{\sigma(i)}\| = \|2\mathbf{q}_i - \mathbf{1}\|.$$

Applying Lemma 2 with $\varepsilon = \frac{(2h_i - 1)\rho_*}{20} < \frac{\|2\mathbf{q}_i - \mathbf{1}\|}{16}$, we obtain an upper bound on the term (a):

$$\begin{aligned}
 \mathbb{P} \left\{ \|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}_{\sigma(i)}\| \geq \frac{\rho_*}{4\|2\mathbf{p}_{\sigma(i)} - 1\|} \right\} &= \mathbb{P} \left\{ \|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}_{\sigma(i)}\| \geq \frac{5}{(2h_i - 1)\|2\mathbf{p}_{\sigma(i)} - 1\|} \frac{(2h_i - 1)\rho_*}{20} \right\} \\
 &= \mathbb{P} \left\{ \|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}_{\sigma(i)}\| \geq \frac{5}{\|2\mathbf{q}_i - 1\|} \frac{(2h_i - 1)\rho_*}{20} \right\} \\
 &\leq 2L \exp \left(-\frac{(2h_i - 1)^2 \rho_*^2 T w}{200 L n} \right) \\
 &= \exp \left(-\frac{(2h_i - 1)^2 \rho_*^2 T w}{200 L n} (1 + o(1)) \right).
 \end{aligned}$$

Step 3. Upper bound of the term (b). Next, we establish the following claim:

$$\sum_{k'=1}^K \mathbb{P} \left\{ \|\xi_{k'} - \tilde{\mathbf{r}}_{k'}\| \geq \frac{\rho_*}{4|2\eta - 1|} \right\} \leq \exp \left(-n \left(\Theta \left(\frac{(T/n)^{\frac{1}{2}}}{\log(T/n)} \right) \right) \right). \tag{18}$$

To this aim, we first show that a large fraction of the items satisfy $\|\hat{\mathbf{r}}^v - \tilde{\mathbf{r}}^v\| \leq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}}$.

Applying Lemma 2 with $\varepsilon = \frac{\|2\mathbf{q}_i - 1\|}{20} \left(\frac{n}{T} \right)^{\frac{1}{4}}$, we get:

$$\begin{aligned}
 \mathbb{P} \left\{ \|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}^i\| \geq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}} \right\} &= \mathbb{P} \left\{ \|\hat{\mathbf{r}}^i - \tilde{\mathbf{r}}^i\| \geq \frac{5}{\|2\mathbf{q}_i - 1\|} \frac{\|2\mathbf{q}_i - 1\|}{20} \left(\frac{n}{T} \right)^{\frac{1}{4}} \right\} \\
 &\leq 2L \exp \left(-\frac{\|2\mathbf{q}_i - 1\|^2 w}{200 L} \left(\frac{T}{n} \right)^{\frac{1}{2}} \right) \\
 &\leq 2L \exp \left(-\frac{h_*^2 \rho_*^2 w}{200L} \left(\frac{T}{n} \right)^{\frac{1}{2}} \right).
 \end{aligned} \tag{19}$$

Define $p_{\max} := \max_{v \in \mathcal{I}_k} \mathbb{P} \left\{ \|\hat{\mathbf{r}}^v - \tilde{\mathbf{r}}^v\| \geq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}} \right\}$. Then from (19),

$p_{\max} \leq \exp \left(-\Theta \left(\left(\frac{T}{n} \right)^{\frac{1}{2}} \right) \right)$. Further define S as the number of the items in \mathcal{I} that satisfy

$\|\hat{\mathbf{r}}^v - \tilde{\mathbf{r}}^v\| \leq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}}$, i.e., $S = \sum_{v \in \mathcal{I}} \mathbb{1}_{\{\|\hat{\mathbf{r}}^v - \tilde{\mathbf{r}}^v\| \leq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}}\}}$. Since $\hat{\mathbf{r}}^1, \dots, \hat{\mathbf{r}}^n$ are independent random variables, $\mathbb{1}_{\{\|\hat{\mathbf{r}}^1 - \tilde{\mathbf{r}}^1\| \leq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}}\}}, \dots, \mathbb{1}_{\{\|\hat{\mathbf{r}}^n - \tilde{\mathbf{r}}^n\| \leq \frac{1}{4} \left(\frac{n}{T} \right)^{\frac{1}{4}}\}}$ are independent Bernoulli random variables. From Chernoff bound, we get:

$$\begin{aligned}
 \mathbb{P}\left\{S \geq \frac{n}{\log\left(\frac{T}{n}\right)}\right\} &\leq \inf_{\lambda>0} \frac{\mathbb{E}[\exp(\lambda S)]}{\exp\left(\lambda \frac{n}{\log\left(\frac{T}{n}\right)}\right)} \leq \inf_{\lambda>0} \frac{(1 + p_{\max}(e^\lambda - 1))^n}{\exp\left(\lambda \frac{n}{\log\left(\frac{T}{n}\right)}\right)} \\
 &\leq \inf_{\lambda>0} \exp\left(np_{\max}(e^\lambda - 1) - \lambda \frac{n}{\log\left(\frac{T}{n}\right)}\right) \\
 &\stackrel{(i)}{\leq} \exp\left(np_{\max}\left(\frac{1}{p_{\max}} - 1\right) - \log \frac{1}{p_{\max}} \frac{n}{\log\left(\frac{T}{n}\right)}\right) \\
 &= \exp\left(-n\left(\Theta\left(\frac{(T/n)^{\frac{1}{2}}}{\log(T/n)}\right) + p_{\max} - 1\right)\right) \\
 &\leq \exp\left(-n\left(\Theta\left(\frac{(T/n)^{\frac{1}{2}}}{\log(T/n)}\right)\right)\right)
 \end{aligned}$$

where for (i), we set $\lambda = \log \frac{1}{p_{\max}}$. Therefore, to prove (18), it suffices to show that:

$$\max_{1 \leq k \leq K} \|\xi_k - \tilde{r}_k\| \leq \frac{\rho_*}{4|2\eta - 1|} \quad \text{when} \quad S \leq \frac{n}{\log\left(\frac{T}{n}\right)}.$$

Assume that $S \leq \frac{n}{\log\left(\frac{T}{n}\right)}$. Then, every v having $\min_{1 \leq k \leq K} \|\hat{r}^v - \tilde{r}_k\| \geq 2\left(\frac{n}{T}\right)^{\frac{1}{4}}$ cannot be a center node (i.e., one the i_k^* for $k = 1, \dots, K$). This is due to the following facts:

- (i) $|T_v| \leq \frac{n}{\log\left(\frac{T}{n}\right)}$ when $\min_{1 \leq k \leq K} \|\hat{r}^v - \tilde{r}_k\| \geq 2\left(\frac{n}{T}\right)^{\frac{1}{4}}$, since for all w such that $\|\hat{r}^w - \tilde{r}^w\| \leq \frac{1}{4}\left(\frac{n}{T}\right)^{\frac{1}{4}}$, $\|\hat{r}^v - \hat{r}^w\| \geq \|\hat{r}^v - \tilde{r}^w\| - \|\tilde{r}^w - \hat{r}^w\| \geq \frac{3}{2}\left(\frac{n}{T}\right)^{\frac{1}{4}}$.
- (ii) $|T_v| \geq \alpha_k n - \frac{n}{\log\left(\frac{T}{n}\right)}$ when $\|\hat{r}^v - \tilde{r}_k\| \leq \frac{1}{2}\left(\frac{n}{T}\right)^{\frac{1}{4}}$, since for all $w \in \mathcal{I}_k$ such that $\|\hat{r}^w - \tilde{r}^w\| \leq \frac{1}{4}\left(\frac{n}{T}\right)^{\frac{1}{4}}$, $\|\hat{r}^v - \hat{r}^w\| \leq \|\hat{r}^v - \tilde{r}_k\| + \|\tilde{r}_k - \hat{r}^w\| \leq \frac{3}{4}\left(\frac{n}{T}\right)^{\frac{1}{4}}$.

Therefore, when $\frac{T}{n} = \omega(1)$,

$$\|\hat{r}_{i_k^*} - \tilde{r}_k\| \leq 2\left(\frac{n}{T}\right)^{\frac{1}{4}}. \tag{20}$$

Let \mathcal{R}_k denote the set of items \mathcal{S}_k before computing ξ_k (\mathcal{S}_k used for the calculation of ξ_k) – see the algorithm. Then, from (20) and the definition of \mathcal{S}_k before computing ξ_k ,

$$\|\hat{r}^v - \tilde{r}_k\| \leq \|\hat{r}^v - \hat{r}_{i_k^*}\| + \|\hat{r}_{i_k^*} - \tilde{r}_k\| \leq 3\left(\frac{n}{T}\right)^{\frac{1}{4}} \quad \text{for all } v \in \mathcal{R}_k.$$

From the above inequality and Jensen’s inequality,

$$\begin{aligned} \|\xi_k - \tilde{r}_k\| &= \left\| \frac{\sum_{v \in \mathcal{R}_k} \hat{r}^v}{|\mathcal{R}_k|} - \tilde{r}_k \right\| \\ &\leq \frac{1}{|\mathcal{R}_k|} \left(\sum_{v \in \mathcal{R}_k} \|\hat{r}^v - \tilde{r}_k\| \right) \leq 3\left(\frac{n}{T}\right)^{\frac{1}{4}}. \end{aligned}$$

Therefore, when $T = \omega(n)$,

$$\max_{1 \leq k \leq K} \|\xi_k - \tilde{r}_k\| \leq \frac{\rho_*}{4|2\eta - 1|} \quad \text{when } S \leq \frac{n}{\log\left(\frac{T}{n}\right)},$$

which concludes the proof of (18).

The proof of theorem is completed by remarking that when $T = o(n^2)$, then

$$\frac{T}{n} = o\left(n \frac{(T/n)^{\frac{1}{2}}}{\log(T/n)}\right).$$

This implies that the upper bound we derived for the term (a) is dominating the upper bound of the term (b). Finally,

$$\mathbb{P}\{i \in \mathcal{E}\} \leq \exp\left(-\frac{(2h_i - 1)^2 \rho_*^2 T w}{200 L n} (1 + o(1))\right).$$

□

4.2 Adaptive selection strategy

Algorithm 2 Adaptive Clustering Algorithm.

Input: K, T
Initialization: $\tau \leftarrow \left\lfloor \frac{T}{4 \log \left(\frac{T}{n} \right)} \right\rfloor$
 $\hat{p}_{k\ell} \leftarrow 0$ for all $\ell \in [L], k \in [K]$
 $\hat{h}_i \leftarrow 0.5$ for all $i \in \mathcal{I}$
while $t < T$ **do**
 $Y_{i\ell} \leftarrow \sum_{t'=1}^t \mathbb{1}[i \in \mathcal{W}_{t'}, \ell_{t'} = \ell]$
 $\hat{q}_{i\ell} \leftarrow \frac{\sum_{t'=1}^t \mathbb{1}[i \in \mathcal{W}_{t'}, \ell_{t'} = \ell, X_{i\ell t'} = 1]}{Y_{i\ell}}$ if $Y_{i\ell} > 0, \frac{1}{2}$ otherwise
if $\text{mod}(t, \tau) = 0$ **then**
 Use Algorithm 1 with input K, t to obtain the estimated clusters $\{\mathcal{S}_k\}_{k=1, \dots, K}$
Estimate statistical parameters:
 $\hat{p}_{k\ell} \leftarrow \frac{1}{2} \left(\frac{\sum_{i \in \mathcal{S}_k} (2\hat{q}_{i\ell} - 1)}{|\mathcal{S}_k|} + 1 \right)$ for all $\ell \in [L], k \in [K]$
 $\hat{h}_i \leftarrow \max \left\{ 0.5, \arg \min_{h' \in [1/2, 1]} \left\{ \sum_{\ell \in [L]} Y_{i\ell} \text{KL}(h' \hat{p}_{\hat{\sigma}(i)\ell} + \bar{h}' \bar{p}_{\hat{\sigma}(i)\ell}, \hat{q}_{i\ell}) \right\} - \sqrt{\frac{\log(t)}{10 \sum_{\ell} Y_{i\ell}}} \right\}$,
 for all $i \in \mathcal{I}$
end if
Adaptive item question selection:
 $i^* \leftarrow \arg \min_{i \in \mathcal{I}} \hat{d}_i$ where $\hat{d}_i = d_i(h'_i, k'_i, \mathbf{Y})$,
 where $k'_i = \arg \min_{k' \neq \hat{\sigma}(i)} \min_{h' \in [1/2, 1]} d_i(h', k', \mathbf{Y})$,
 $h'_i = \arg \min_{h' \in [1/2, 1]} d_i(h', k'_i, \mathbf{Y})$,
 and $d_i(h', k', \mathbf{Y}) = \sum_{\ell=1}^L Y_{i\ell} \text{KL}(h' \hat{p}_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, \hat{h}_i \hat{p}_{\hat{\sigma}(i)\ell} + \bar{h}_i \bar{p}_{\hat{\sigma}(i)\ell})$.
 $\ell^* \leftarrow \arg \max_{\ell} \min_{k' \neq \hat{\sigma}(i^*)} \text{KL}(h'_{i^*} \hat{p}_{k'\ell} + h'_{i^*} \bar{p}_{k'\ell}, \hat{h}_{i^*} \hat{p}_{\hat{\sigma}(i^*)\ell} + \bar{h}_{i^*} \bar{p}_{\hat{\sigma}(i^*)\ell})$.
 Present items \mathcal{W}_t (including i^*) with the smallest \hat{d}_i with question ℓ^* .
end while
Output: $\{\mathcal{S}_k\}_{k \in [K]}$

Our adaptive (item, question) selection and clustering algorithm is described in Algorithm 2. The design of the adaptive (item, question) selection strategy is inspired by the derivation of the information-theoretical error lower bounds. The algorithm maintains estimates of the model parameters \mathbf{p} and \mathbf{h} and of the clusters $\{\mathcal{I}_k\}_{k \in [K]}$. These estimates, denoted by $\hat{\mathbf{p}}, \hat{\mathbf{h}}$, and $\{\mathcal{S}_k\}_{k \in [K]}$, respectively, are updated every $\tau = T/(4 \log(T/n))$ users. More precisely, we use Algorithm 1 to compute $\{\mathcal{S}_k\}_{k \in [K]}$, and from there, we update the estimates as:

$$\hat{p}_{k\ell} = \frac{1}{2} \left(\frac{\sum_{i \in \mathcal{S}_k} (2\hat{q}_{i\ell} - 1)}{|\mathcal{S}_k|} + 1 \right),$$

$$\hat{h}_i = \arg \min_{h' \in [1/2, 1]} \left\{ \sum_{\ell \in [L]} Y_{i\ell} \text{KL}(h' \hat{p}_{\hat{\sigma}(i)\ell} + \bar{h}' \bar{p}_{\hat{\sigma}(i)\ell}, \hat{q}_{i\ell}) \right\},$$

where $Y_{i\ell}$ is the number of times where question ℓ has been asked for item i so far, and where $\hat{\sigma}(i)$ corresponds to the estimated cluster of i (i.e., $i \in \mathcal{S}_{\hat{\sigma}(i)}$). Let $\mathbf{Y} := (Y_{i\ell})_{i \in \mathcal{I}, \ell \in [L]}$.

Now using the same arguments as those used to derive error lower bounds, we may estimate that after seeing the t -th user, a lower bound on the mis-classification error for item i is $\exp(-\hat{d}_i(\mathbf{Y}))$, where

$$\hat{d}_i(\mathbf{Y}) := \min_{k' \neq \hat{\sigma}(i)} \min_{h' \in [1/2, 1]} \sum_{\ell=1}^L Y_{i\ell} \text{KL}(h' \hat{p}_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, \hat{h}_i \hat{p}_{\hat{\sigma}(i)\ell} + \bar{\hat{h}}_i \bar{p}_{\hat{\sigma}(i)\ell}).$$

The above lower bounds are heuristic in nature, as they are based solely on estimated parameters and clusters. These are derived from the divergence $\mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})$ using (5), with a particular emphasis on the adjustable parameters for item i . This approach takes a pessimistic view of the hardness parameters, with the exception of those for item i . Revisiting the scenario of Example 1, there is only one question ($L = 1$) and the adaptability of the algorithm is principally determined by how the budget T is allocated among the items. Observe that, when h_i is estimated to be small, the value of $\text{KL}(h' \hat{p}_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, \hat{h}_i \hat{p}_{\hat{\sigma}(i)\ell} + \bar{\hat{h}}_i \bar{p}_{\hat{\sigma}(i)\ell})$ tends to be small. Conversely, when h_i is estimated to be large, the value of $\text{KL}(h' \hat{p}_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, \hat{h}_i \hat{p}_{\hat{\sigma}(i)\ell} + \bar{\hat{h}}_i \bar{p}_{\hat{\sigma}(i)\ell})$ tends to be large. Therefore, the more difficult the item i is, the greater the need for a larger Y_{i1} , and the higher the frequency of it being selected. Analyzing the accuracy of these lower bounds is particularly challenging (it is hard to analyze the estimated item hardness \hat{h}_i). Using these estimated lower bounds, we select the items and the question to be asked next. We put in the list \mathcal{W}_t the w items with the smallest $\hat{d}_i(\mathbf{Y})$. The question ℓ is chosen to maximize the term: $\min_{k' \neq \hat{\sigma}(i^*)} \text{KL}(h'_{i^*} \hat{p}_{k'\ell} + h'_{i^*} \bar{p}_{k'\ell}, \hat{h}_{i^*} \hat{p}_{\hat{\sigma}(i^*)\ell} + \bar{\hat{h}}_{i^*} \bar{p}_{\hat{\sigma}(i^*)\ell})$, where $i^* = \arg \min_{i \in \mathcal{I}} \hat{d}_i(\mathbf{Y})$ (see

Algorithm 2 for the details). Note that the question is selected by considering the item i^* that seems to be the most difficult to classify.

Note that in order to reduce the computational complexity of the algorithm, we may replace the KL function in the definition of d_i by a simple quadratic function (as suggested in the proof of Proposition 1). This simplifies the minimization problem over h' to find h'_i . We actually have an explicit expression for h'_i with this modification.

The computational complexity of the adaptive algorithm (Algorithm 2 in Appendix) is: $\mathcal{O}(n^2 T / \tau) = \mathcal{O}(n^2 \log(T/n))$. As in the uniform case, by choosing a small ($\log n$) subset of items (and not all the items in \mathcal{I}) to compute centroids (T_i), one can reduce the computational complexity to: $\mathcal{O}(n \log(n) \log(T/n))$. We provide experimental evidence on the superiority of our adaptive algorithm in the following sections.

5 Numerical experiments: synthetic data

In this section, we evaluate the performance of our algorithms on synthetic data. We consider different models. the problem investigated here is different from those one may find in the crowdsourcing or Stochastic Block Model literature. Hence, we cannot compare our algorithms to existing algorithms developed in this literature. Instead we focus on comparing the performance of our nonadaptive and adaptive algorithms.

Model 1: heterogeneous items with dummy questions: Consider $n = 1000$ items and two clusters ($K = 2$) of equal sizes. The hardness of the items are i.i.d., picked uniformly at random in the interval $[0.55, 1]$. We ask each user to answer one of four questions. The answers' statistics are as follows: for cluster $k = 1$, $\mathbf{p}_1 = (0.01, 0.99, 0.5, 0.5)$ and for cluster $k = 2$, $\mathbf{p}_2 = (0.99, 0.01, 0.5, 0.5)$. Note that only half of the questions ($\ell = 1, 2$) are useful; the other questions ($\ell = 3, 4$) generate completely random answers for both clusters. Figure 2 (top) plots the error rate averaged over all items and over 100 instances of our

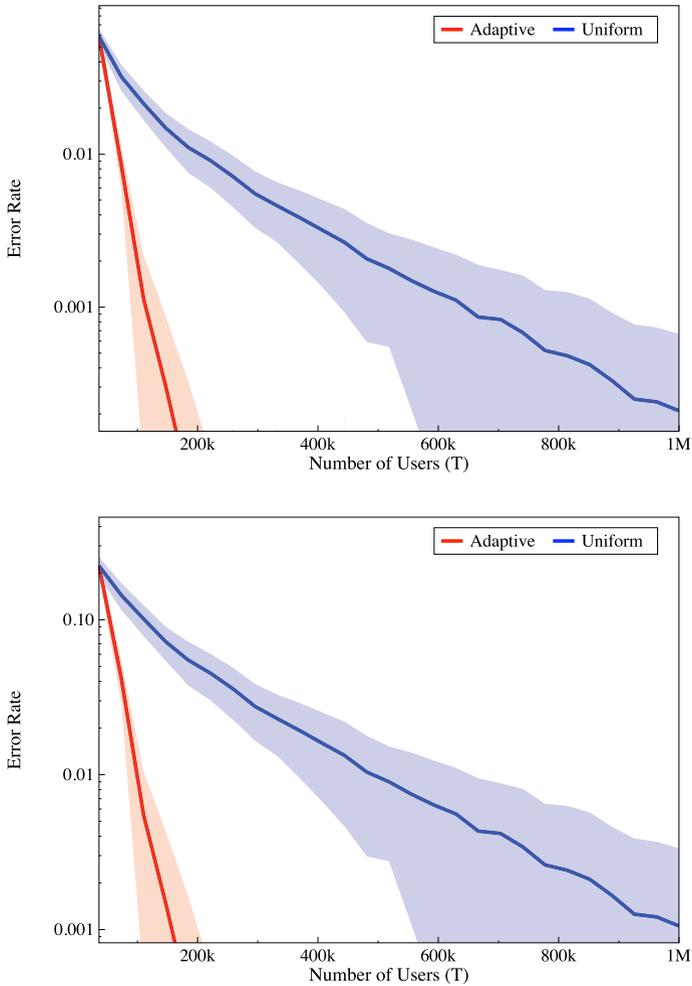


Fig. 2 Model 1: (top) Global error rate vs. number of users. (bottom) Error rate for the 20% most difficult items vs. number of users. One standard deviations are shown using shaded areas

algorithms. Under both algorithms, the error rate decays exponentially with the budget T , as expected from the analysis. Selecting items and questions in an adaptive manner brings significant performance improvements. For example, after collecting the answers from $t = 200k$, the adaptive algorithm recovers the clusters exactly for most of the instances, whereas the algorithm using uniform selection does not achieve exact recovery even with $t = 1000k$ users. In particular, the adaptive algorithm is able to reduce the error rates on the 20% most difficult items, i.e., items that have the 20% smallest h_i . In Fig. 2 (bottom), we present the error rate of these items. The error rates for these most difficult items are significantly reduced by being adaptive. In Fig. 3, we present the evolution over time of the

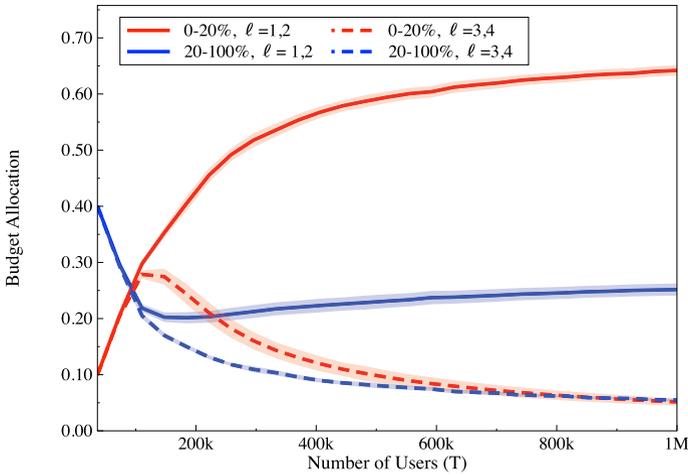


Fig. 3 Model 1: The budget allocation under the adaptive algorithm vs. number of users. Items and questions are grouped into 4 categories, e.g. (0 – 20%, $\ell = 1, 2$) is the category regrouping the 20% most difficult items and questions $\ell = 1, 2$. One standard deviations are shown using shaded areas

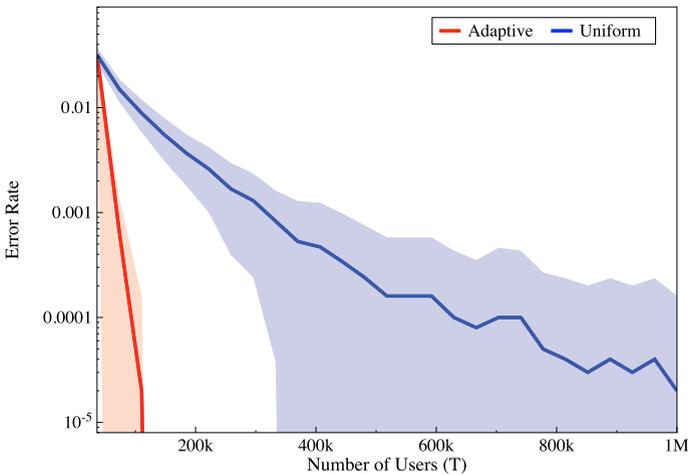


Fig. 4 Model 2: Global error rate vs. number of users. One standard deviations are shown using shaded areas

budget allocation observed under our adaptive algorithm. We group items and questions into four categories. For example, one category corresponds to the question $\ell = 1, 2$ and to the 20% most difficult items. As expected, the adaptive algorithm learns to select relevant questions ($\ell = 1, 2$) with hard items more and more often as time evolves.

Model 2: heterogeneous items without dummy questions. This model is similar to Model 1, except that we remove the dummy questions $\ell = 3, 4$, i.e., we set $p_1 = (0.01, 0.99)$ and $p_2 = (0.99, 0.01)$. The performance of our algorithms are shown in

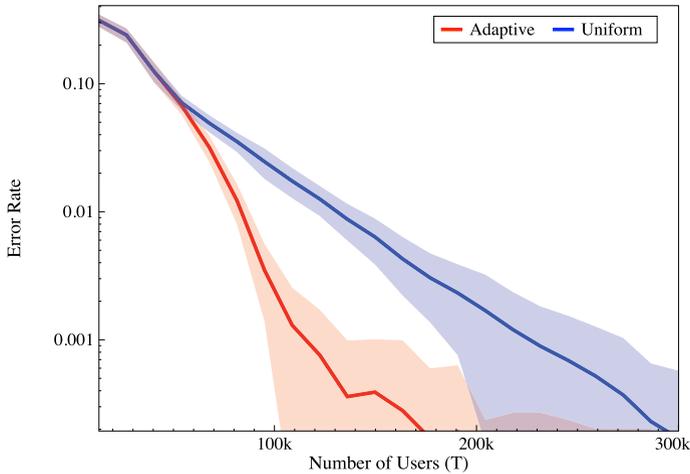


Fig. 5 Model 3: Global error rate vs. number of users. One standard deviations are shown using shaded areas

Fig. 4. Overall, compared to Model 1, the error rates are better. For example, exact cluster recovery is achieved using only 100k users for almost all instances.

Model 3: homogeneous items with dummy questions. Here we study the homogeneous scenario where all items have the same hardness: $h_i = 1, \forall i \in \mathcal{I}$. We still have 1000 items grouped into two clusters of equal sizes. We set $\mathbf{p}_1 = (0.3, 0.2, 0.2, 0.2)$, $\mathbf{p}_2 = (0.7, 0.2, 0.2, 0.2)$ (question $\ell = 2, 3, 4$ are useless). The performance of the algorithms is shown in Fig. 5. The adaptive algorithm exhibits better error rates than the algorithm with uniform selection, although the improvement is not as spectacular as in heterogeneous models where adaptive algorithms can gather more information about difficult items. In homogeneous models, the adaptive algorithm remains better because it selects questions wisely.

6 Numerical experiments: real-world data

Finally, we use real-world data to assess the performance of our algorithms. Finding data that would fit our setting exactly (e.g. several possible questions) is not easy. We restrict our attention here to scenarios with a single question, but with items with different hardnesses. We use the waterbird dataset by Welinder et al. (2010). This dataset contains

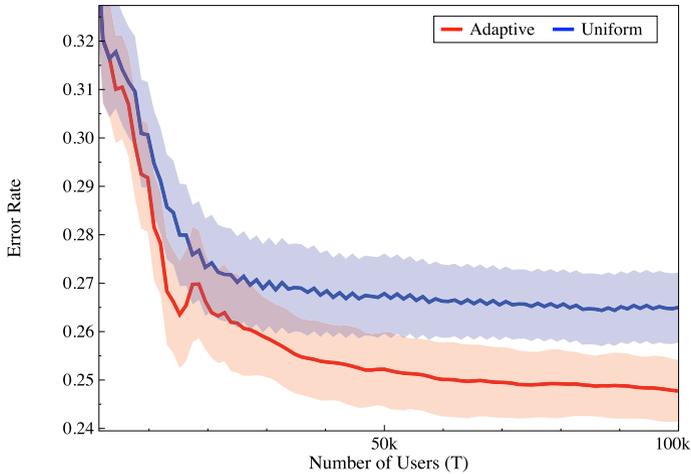


Fig. 6 Real-World data: global error rate vs. number of users. One standard deviations are shown using shaded areas

50 images of Mallards (a kind of duck) and 50 images of Canadian Goose (not a duck). The dataset reports the feedback of 40 users per image, collected using Amazon Mturk: each user is asked whether the image is that of a duck. This scenario mirrors the one outlined in Example 1 in Introduction. Each image is unique in the sense that the orientation of the animal varies, the brightness and contrasts are different, etc. We hence have a good heterogeneity in terms of item hardness. Actually, the classification task is rather difficult, and the users' answers seem very noisy – overall answers are correct 76% of the time.

From this small dataset, we generated a larger dataset containing 1000 images (by just replicating images). To emulate the sequential nature of our clustering problem, in each round, we pick a user uniformly at random (with replacement), and observe her answers to the selected images.

The error rates of both algorithms are shown in Fig. 6. The global error rate is averaged over 100 instances. Both algorithms have rather low performance, which can be explained by the inherent hardness of the learning task. The adaptive algorithm becomes significantly better after $t = 20k$ users. this can be explained as follows. The adaptive algorithm needs to estimate the hardness of items before being efficient. Until the algorithm gathers enough answers on item i , its estimate of \hat{h}_i remains close to 0.5. As a consequence, the algorithm continues to pick items uniformly at random. As soon as the algorithm gets better estimates of the items' hardnesses, it starts selecting items with strong preferences.

7 Conclusion

In this paper, we analyzed the problem of clustering complex items using very simple binary feedback provided by users. A key aspect of our problem is that it takes into account the fact that some items are inherently more difficult to cluster than some others. Accounting for this heterogeneity is critical to get realistic models, and is unfortunately

not investigated often in the literature on clustering and community detection (e.g. that on Stochastic Block Model). The item heterogeneity also significantly complicates any theoretical development.

For the case where data is collected uniformly (each item receives the same amount of user feedback), we derived a lower bound of the clustering error rate for any individual item, and we developed a clustering algorithm approaching the optimal error rate. We also investigated adaptive algorithms, under which the user feedback is received sequentially, and can be adapted to past observations. Being adaptive allows to gather more feedback for more difficult items. We derived a lower bound of the error rate that holds for any adaptive algorithm. Based on our lower bounds, we devised an adaptive algorithm that smartly select items and the nature of the feedback to be collected. We evaluated our algorithms on both synthetic and real-world data. These numerical experiments support our theoretical results, and demonstrate that being adaptive leads to drastic performance gains.

Appendix A: Table of notations

Problem-specific notations

\mathcal{I}	Set of items
\mathcal{I}_k	Set of items in the item cluster k
n	Number of items
K	Number of item clusters
$\sigma(i)$	Cluster index of item i
$\alpha := (\alpha_1, \dots, \alpha_K)$	Fractions of items in each cluster
w	The number of items presented at the same time
\mathcal{W}_t	Set of items presented to the t -th user
L	Number of possible questions
ℓ_t	Question asked to the t -th user
T	Total number of user arrivals within the time horizon
$\mathbf{p} := (p_{k\ell})_{k \in [K], \ell \in [L]}$	Statistical parameterization of items in cluster k for the question ℓ
$\mathbf{h} := (h_i)_{i \in \mathcal{I}}$	Hardness parameter for item i
\mathcal{M}	Statistical models parameterized by (\mathbf{p}, \mathbf{h})
$X_{i\ell t}$	Binary feedback from t -th user for item i and for question ℓ
$q_{i\ell} := h_i p_{k\ell} - \bar{h}_i \bar{p}_{k\ell}$	Probability of positive answer to item i for question ℓ
h_*	Minimum hardness across items, see Assumption (A1)
ρ_*	Minimum separation between different clusters, see Assumption (A1)
η	Homogeneity parameter among clusters, see Assumption (A2)
Ω	Set of all models satisfying Assumptions (A1) and (A2)
$r_{k\ell}$	Value of $2p_{k\ell} - 1$
\mathcal{E}^π	Set of misclassified items by the algorithm π
$\varepsilon_i^\pi(n, T)$	Probability that the item i is misclassified after the T -th user arrived under the algorithm π
$\varepsilon^\pi(n, T)$	Expected proportion of misclassified items after the T -th user arrived under the algorithm π

Problem-specific notations

$Y_{i\ell}$	Number of times the item i is presented together with the question ℓ
$y_{i\ell}$	Normalized expected number of times the question ℓ is asked for the item i under some fixed algorithm
$\mathcal{D}_{\mathcal{M}}^U(i)$	Divergence for the misclassification of item i with the model \mathcal{M} under uniform item selection strategy
$\tilde{\mathcal{D}}_{\mathcal{M}}^U$	Global divergence with the model \mathcal{M} under uniform item selection strategy
$\mathcal{D}_{\mathcal{M}}^A(i, \mathbf{y})$	Divergence for the misclassification of item i with the model \mathcal{M} under some adaptive item selection strategy satisfying $\mathbb{E}[Y_{i\ell}] = \frac{T_w}{n} y_{i\ell}$
$\tilde{\mathcal{D}}_{\mathcal{M}}^A$	Global divergence with the model \mathcal{M} under the optimal adaptive item selection strategy

Generic notations

\hat{a}	Estimated value of a
$[a]$	Set of positive integers upto a , i.e., $\{1, \dots, a\}$
$\mathbb{1}\{A\}$	Indicator function: 0 when A is false, 1 when A is true
$\ \mathbf{x}\ $	ℓ_∞ norm of \mathbf{x} , i.e., $\ \mathbf{x}\ = \max_i x_i$
$\ \mathbf{x}\ _2$	ℓ_2 norm of \mathbf{x}
\bar{a}	Value of $1 - a$
$\mathbb{P}(A)$	Probability that event A occurs
$\mathbb{E}[a]$	Expected value of a
$\text{KL}(a, b)$	Kullback–Leibler divergence between Bernoulli distributions with means a and b

Appendix B: Proof of Proposition 1

For given $i \in \mathcal{I}$, let $k = \sigma(i)$ and $k' \in [K]$ be such that:

$$\mathcal{D}_{\mathcal{M}}^U(i) = \min_{h' \in \{(h_*+1)/2, 1\}} \frac{1}{L} \sum_{\ell=1}^L \text{KL}(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, q_{i\ell}).$$

Upper bound. Recalling the definition of $q_{i\ell} := h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}$, it follows that for any $h' \in [(h_* + 1)/2, 1]$,

$$\begin{aligned} \mathcal{D}_{\mathcal{M}}^U(i) &\leq \frac{1}{L} \sum_{\ell=1}^L \text{KL}(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}, h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}) \\ &\leq \frac{1}{L} \sum_{\ell=1}^L ((h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}) - (h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}))^2 \left(\frac{1}{q_{i\ell}} + \frac{1}{\bar{q}_{i\ell}} \right) \\ &\leq \frac{2}{L\eta} \sum_{\ell=1}^L ((h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}) - (h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell}))^2 \\ &= \frac{2}{L\eta} \sum_{\ell=1}^L \left(\frac{(2h' - 1)(2p_{k'\ell} - 1) + 1}{2} - \frac{(2h_i - 1)(2p_{k\ell} - 1) + 1}{2} \right)^2 \\ &= \frac{1}{2L\eta} (2h_i - 1)^2 \left\| \frac{2h' - 1}{2h_i - 1} \mathbf{r}_{k'} - \mathbf{r}_k \right\|_2^2. \end{aligned}$$

where the second inequality is from the comparison between the KL divergence from χ^2 -divergence and the third inequality is from (A2), i.e., $q_{i\ell} \in [\eta, 1 - \eta]$. Now observing that $h' \in [(h_* + 1)/2, 1]$ implies $\frac{h_*}{2h_i - 1} \leq \frac{2h' - 1}{2h_i - 1} \leq \frac{1}{2h_i - 1}$. Taking the minimum over $h' \in [(h_* + 1)/2, 1]$, we obtain the upper bound.

Lower bound. Using Pinsker’s inequality, we obtain:

$$\begin{aligned} \mathcal{D}_{\mathcal{M}}^U(i) &\geq \frac{2}{L} \sum_{\ell=1}^L |(h' p_{k'\ell} + \bar{h}' \bar{p}_{k'\ell}) - (h_i p_{k\ell} + \bar{h}_i \bar{p}_{k\ell})|^2 \\ &= \frac{1}{2L} \sum_{\ell=1}^L |(2h' - 1)(2p_{k'\ell} - 1) - (2h_i - 1)(2p_{k\ell} - 1)|^2 \\ &= \frac{1}{2L} \|(2h' - 1)\mathbf{r}_{k'} - (2h_i - 1)\mathbf{r}_k\|_2^2 \\ &= \frac{(2h_i - 1)^2}{2L} \left\| \left(\frac{2h' - 1}{2h_i - 1} \right) \mathbf{r}_{k'} - \mathbf{r}_k \right\|_2^2 \\ &\geq \frac{(2h_i - 1)^2}{2L} \min_{\frac{h_*}{2h_i - 1} \leq \frac{2h' - 1}{2h_i - 1} \leq \frac{1}{2h_i - 1}} \|\mathbf{c}_{k'} - \mathbf{r}_k\|_2^2, \end{aligned}$$

where for the last inequality, we again use the fact that $h' \in [(h_* + 1)/2, 1]$ implies $\frac{h_*}{2h_i - 1} \leq \frac{2h' - 1}{2h_i - 1} \leq \frac{1}{2h_i - 1}$. This completes the proof of Proposition 1. Note that we can further write:

$$\begin{aligned} \mathcal{D}_{\mathcal{M}}^U(i) &\geq \frac{1}{2L} (2h_i - 1)^2 \min_{\frac{h_*}{2h_i - 1} \leq \frac{2h' - 1}{2h_i - 1} \leq \frac{1}{2h_i - 1}} \|\mathbf{c}_{k'} - \mathbf{r}_k\|_2^2 \\ &\geq \frac{1}{2L} (2h_i - 1)^2 \rho_*^2, \end{aligned}$$

using the relationship between the ℓ_∞ -norm and the Euclidean norm and (A1). □

Appendix C: Proof of Lemma 1

$\mathbb{E}_{\mathcal{N}}[L]$ can be obtained as follows:

$$\begin{aligned}
 \mathbb{E}_{\mathcal{N}}[\mathcal{L}] &= \sum_{\ell=1}^L \sum_{t=1}^T \mathbb{1}[i \in \mathcal{W}_t \text{ and } \ell_t = \ell] \mathbb{E}_{\mathcal{N}} \left[\mathbb{1}[X_{i\ell_t} = +1] \log \frac{q'_{\ell}}{q_{i\ell}} + \mathbb{1}[X_{i\ell_t} = -1] \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right] \\
 &= \sum_{\ell=1}^L \frac{Tw}{Ln} \left(q'_{\ell} \log \frac{q'_{\ell}}{q_{i\ell}} + \bar{q}'_{\ell} \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right) \\
 &= \frac{Tw}{Ln} \sum_{\ell=1}^L \text{KL}(q'_{\ell}, q_{i\ell}) \\
 &= \frac{Tw}{n} \mathcal{D}_{\mathcal{M}}^U(i).
 \end{aligned} \tag{21}$$

To bound the variance of \mathcal{L} , we first decompose \mathcal{L}^2 as follows:

$$\mathcal{L}^2 = \sum_{t=1}^T \mathcal{L}_t^2 + \sum_{t \neq t'} \mathcal{L}_t \mathcal{L}_{t'},$$

where $\mathcal{L}_t := \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell_t} = +1] \log \frac{q'_{\ell}}{q_{i\ell}} + \mathbb{1}[X_{i\ell_t} = -1] \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right)$. We compute \mathcal{L}_t^2 as follows:

$$\begin{aligned}
 \mathcal{L}_t^2 &= \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell_t} = +1] \log \frac{q'_{\ell}}{q_{i\ell}} + \mathbb{1}[X_{i\ell_t} = -1] \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right)^2 \\
 &= \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell_t} = +1] \left(\log \frac{q'_{\ell}}{q_{i\ell}} \right)^2 + \mathbb{1}[X_{i\ell_t} = -1] \left(\log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right)^2 \right) \\
 &\leq \log(1/\eta) \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell_t} = +1] \left| \log \frac{q'_{\ell}}{q_{i\ell}} \right| + \mathbb{1}[X_{i\ell_t} = -1] \left| \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right| \right),
 \end{aligned} \tag{22}$$

where the last inequality follows from the fact that $q_{i\ell} \in [\eta, 1 - \eta]$ under (A2), i.e., $\log \frac{q'_{\ell}}{q_{i\ell}} \leq \log \frac{1}{\eta}$ and $\log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \leq \log \frac{1}{\eta}$. We deduce that:

$$\begin{aligned}
 \mathbb{E}_{\mathcal{N}} \left[\sum_{t=1}^T \mathcal{L}_t^2 \right] &= \sum_{t=1}^T \mathbb{E}_{\mathcal{N}}[\mathcal{L}_t^2] \\
 &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{N}} \left[\log(1/\eta) \mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell_t} = +1] \left| \log \frac{q'_{\ell}}{q_{i\ell}} \right| + \mathbb{1}[X_{i\ell_t} = -1] \left| \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right| \right) \right] \\
 &= \log(1/\eta) \sum_{\ell=1}^L \sum_{t=1}^T \mathbb{1}[i \in \mathcal{W}_t \text{ and } \ell_t = \ell] \mathbb{E}_{\mathcal{N}} \left[\mathbb{1}[X_{i\ell_t} = +1] \left| \log \frac{q'_{\ell}}{q_{i\ell}} \right| + \mathbb{1}[X_{i\ell_t} = -1] \left| \log \frac{\bar{q}'_{\ell}}{\bar{q}_{i\ell}} \right| \right] \\
 &= \log(1/\eta) \frac{Tw}{Ln} \sum_{\ell=1}^L \left(\text{KL}(q_{v\ell}, q_{i\ell}) + \sqrt{\text{KL}(q_{v\ell}, q_{i\ell})} \right),
 \end{aligned} \tag{23}$$

where for the last inequality, we used the Pinsker’s inequality. Moreover we can compute the expectation of $\sum_{t \neq t'} \mathcal{L}_t \mathcal{L}_{t'}$ as follows:

$$\begin{aligned}
 \mathbb{E}_{\mathcal{N}} \left[\sum_{t \neq t'} \mathcal{L}_t \mathcal{L}_{t'} \right] &= \sum_{t \neq t'} \mathbb{E}_{\mathcal{N}} [\mathcal{L}_t \mathcal{L}_{t'}] \\
 &= \sum_{t \neq t'} \mathbb{E}_{\mathcal{N}} \left[\mathbb{1}[i \in \mathcal{W}_t] \sum_{\ell=1}^L \mathbb{1}[\ell_t = \ell] \left(\mathbb{1}[X_{i\ell t} = +1] \log \frac{q'_\ell}{q_{i\ell}} + \mathbb{1}[X_{i\ell t} = -1] \log \frac{\bar{q}'_\ell}{\bar{q}_{i\ell}} \right) \right] \\
 &\cdot \mathbb{E}_{\mathcal{N}} \left[\mathbb{1}[i \in \mathcal{W}_{t'}] \sum_{\ell'=1}^L \mathbb{1}[\ell_{t'} = \ell'] \left(\mathbb{1}[X_{i\ell' t'} = +1] \log \frac{q'_{\ell'}}{q_{i\ell'}} + \mathbb{1}[X_{i\ell' t'} = -1] \log \frac{\bar{q}'_{\ell'}}{\bar{q}_{i\ell'}} \right) \right] \\
 &= \sum_{t \neq t'} \mathbb{1}[i \in \mathcal{W}_t \text{ and } i \in \mathcal{W}_{t'}] \sum_{\ell=1}^L \sum_{\ell'=1}^L \mathbb{1}[\ell_t = \ell \text{ and } \ell_{t'} = \ell'] \text{KL}(q'_\ell, q_{i\ell}) \text{KL}(q'_{\ell'}, q_{i\ell'}) \\
 &= \sum_{\ell=1}^L \sum_{\ell'=1}^L T(T-1) \left(\frac{w}{Ln} \right)^2 \text{KL}(q'_\ell, q_{i\ell}) \text{KL}(q'_{\ell'}, q_{i\ell'}) \\
 &= T(T-1) \left(\frac{w}{Ln} \right)^2 \left(\sum_{\ell=1}^L \text{KL}(q'_\ell, q_{i\ell}) \right) \left(\sum_{\ell'=1}^L \text{KL}(q'_{\ell'}, q_{i\ell'}) \right) \\
 &= T(T-1) \left(\frac{w}{Ln} \right)^2 \left(\sum_{\ell=1}^L \text{KL}(q'_\ell, q_{i\ell}) \right)^2 \\
 &= \mathbb{E}_{\mathcal{N}}[\mathcal{L}]^2 - T \left(\frac{w}{Ln} \right)^2 \left(\sum_{\ell=1}^L \text{KL}(q'_\ell, q_{i\ell}) \right)^2,
 \end{aligned}$$

where for the last equality, we use the expression (21) of $\mathbb{E}_{\mathcal{N}}[\mathcal{L}]$. Combining (23) with the above, it follows that:

$$\mathbb{E}_{\mathcal{N}}[(\mathcal{L} - \mathbb{E}_{\mathcal{N}}[\mathcal{L}])^2] = \mathbb{E}_{\mathcal{N}}[\mathcal{L}^2] - \mathbb{E}_{\mathcal{N}}[\mathcal{L}]^2 \leq \frac{Tw}{Ln} \log(1/\eta) \sum_{\ell=1}^L \text{KL}(q'_\ell, q_{i\ell}) + \sqrt{\text{KL}(q_{i\ell}, q_{i\ell})}.$$

□

Appendix D: Proof of Corollary 1

We have:

$$\begin{aligned}
 \mathcal{D}_{\mathcal{M}}^U(i) &= \min_{h' \in ((h_i+1)/2, 1]} \text{KL}(h'p_{21} + \bar{h}'\bar{p}_{21}, q_{i1}) \\
 &\leq \text{KL}(h_i p_{21} + \bar{h}_i \bar{p}_{21}, q_{i1}) \\
 &\stackrel{(a)}{\leq} \frac{((h_i p_{21} + \bar{h}_i \bar{p}_{21}) - q_{i1})^2}{q_{i1}(1 - q_{i1})} \\
 &\stackrel{(b)}{\leq} \frac{2}{\eta} \left((h_i p_{11} + \bar{h}_i \bar{p}_{11}) - (h_i p_{21} + \bar{h}_i \bar{p}_{21}) \right)^2 \\
 &= \frac{2}{\eta} \left(\frac{(2h_i - 1)(2p_{11} - 1) + 1}{2} - \frac{(2h_i - 1)(2p_{21} - 1) + 1}{2} \right)^2 \\
 &= \frac{2}{\eta} (2h_i - 1)^2 (p_{11} - p_{21})^2,
 \end{aligned}$$

where (a) stems from the relationship between the KL divergence and the χ^2 -divergence and (b) is from (A2). Combining this inequality and Theorem 1 yield Corollary 1. \square

Proof of Lemma 2

We use Hoeffding’s inequality to establish the lemma.

Theorem 4 (Hoeffding’s inequality for bounded independent random variables (Theorem 1 of Hoeffding (1963))) *Let X_1, \dots, X_n be independent random variables with values in $[0, 1]$.*

Denote $\mu = \mathbb{E} \left[\frac{1}{n} \sum_{i=1}^n X_i \right]$. Then, for any $t \geq 0$,

$$\mathbb{P} \left\{ \frac{1}{n} \sum_{i=1}^n X_i - \mu \geq t \right\} \leq \exp(-2nt^2).$$

Lemma 3 *Recall that by definition, $\tau := \lfloor \frac{T_w}{Ln} \rfloor$. For any $\epsilon > 0$, $\|\hat{q}_i - q_i\| \leq \epsilon$ with probability at least $1 - 2L \exp(-2\tau\epsilon^2)$.*

Proof of Lemma 3 Note that the number of times question ℓ is asked for item i is τ . Using Hoeffding’s inequality (Theorem 4), it is straightforward to check: for any $\epsilon > 0$ and $\ell \in [L]$,

$$\mathbb{P} \left\{ \left| \frac{1}{\tau} \sum_{t=1}^{\tau} (X_{i\ell t} - q_{i\ell}) \right| = |\hat{q}_{i\ell} - q_{i\ell}| \geq \epsilon \right\} \leq 2 \exp(-2\tau\epsilon^2).$$

We conclude the proof using the union bound as follows:

$$\mathbb{P} \{ \|\hat{q}_i - q_i\| \geq \epsilon \} \leq \sum_{\ell \in [L]} \mathbb{P} \{ |\hat{q}_{i\ell} - q_{i\ell}| \geq \epsilon \} \leq 2L \exp(-2\tau\epsilon^2).$$

\square

Proof of Lemma 2 By Lemma 3, we have $\|\hat{q}_i - q_i\| \leq \varepsilon$ with probability at least $1 - 2L \exp(-2\varepsilon^2\tau)$. Suppose $\|\hat{q}_i - q_i\| \leq \varepsilon$ and $0 < \varepsilon \leq \frac{\|2q_i - 1\|}{16}$. Using the triangle inequality and the reverse triangle inequality, we have

$$\|2q_i - 1\| - 2\|q_i - \hat{q}_i\| \leq \|2\hat{q}_i - 1\| \leq \|2q_i - 1\| + 2\|q_i - \hat{q}_i\|,$$

Therefore,

$$\|2q_i - 1\| - 2\varepsilon \leq \|2\hat{q}_i - 1\| \leq \|2q_i - 1\| + 2\varepsilon$$

which implies that:

$$\frac{1}{\|2q_i - 1\|} \frac{1}{1 + \frac{2\varepsilon}{\|2q_i - 1\|}} \leq \frac{1}{\|2\hat{q}_i - 1\|} \leq \frac{1}{\|2q_i - 1\|} \frac{1}{1 - \frac{2\varepsilon}{\|2q_i - 1\|}}.$$

From $0 < \varepsilon \leq \frac{\|2q_i - 1\|}{16}$, we have $0 < \frac{2\varepsilon}{\|2q_i - 1\|} \leq \frac{1}{8}$. Now observe that we have:

$$1 - x \leq \frac{1}{1 + x} \quad \text{and} \quad \frac{1}{1 - x} \leq 1 + \frac{8}{7}x,$$

for all x such that $0 < x < \frac{1}{8}$. Then we obtain:

$$\frac{1}{\|2q_i - 1\|} \left(1 - \frac{2\varepsilon}{\|2q_i - 1\|}\right) \leq \frac{1}{\|2\hat{q}_i - 1\|} \leq \frac{1}{\|2q_i - 1\|} \left(1 + \frac{8}{7} \frac{2\varepsilon}{\|2q_i - 1\|}\right).$$

Then, there exists $x \in \left[-\frac{2\varepsilon}{\|2q_i - 1\|}, \frac{16\varepsilon}{7\|2q_i - 1\|}\right]$ such that $\frac{1}{\|2\hat{q}_i - 1\|} = \frac{1}{\|2q_i - 1\|}(1 + x)$. Using this x , we get:

$$\begin{aligned} \|\hat{r}^i - \tilde{r}^i\| &= \left\| \frac{2\hat{q}_i - 1}{\|2\hat{q}_i - 1\|} - \frac{2q_i - 1}{\|2q_i - 1\|} \right\| \\ &= \frac{1}{\|2q_i - 1\|} \left\| \frac{\|2q_i - 1\|}{\|2\hat{q}_i - 1\|} (2\hat{q}_i - 1) - (2q_i - 1) \right\| \\ &= \frac{1}{\|2q_i - 1\|} \|(1 + x)(2\hat{q}_i - 1) - 2q_i + 1\| \\ &\leq \frac{1}{\|2q_i - 1\|} (\|2(\hat{q}_i - q_i)\| + \|x(2\hat{q}_i - 1)\|) \\ &\leq \frac{1}{\|2q_i - 1\|} (\|2(\hat{q}_i - q_i)\| + \|x(2\hat{q}_i - 2q_i)\| + \|x(2q_i - 1)\|) \\ &\leq \frac{1}{\|2q_i - 1\|} (2\varepsilon + \frac{16}{7}\varepsilon \cdot 2\varepsilon + \frac{16}{7}\varepsilon) \\ &= \frac{1}{\|2q_i - 1\|} \left(\frac{30}{7}\varepsilon + \frac{32}{7}\varepsilon^2\right) \\ &\stackrel{(a)}{\leq} \frac{5\varepsilon}{\|2q_i - 1\|}, \end{aligned}$$

with probability at least $1 - 2L \exp(-2\varepsilon^2\tau)$ for all ε such that $0 < \varepsilon \leq \frac{\|2q_i - 1\|}{16}$, where in (a), we use $5x \geq \frac{30}{7}x + \frac{32}{7}x^2$ for all $0 < x < \frac{1}{16}$ and $0 < \varepsilon \leq \frac{\|2q_i - 1\|}{16} \leq \frac{1}{16}$. This concludes the proof. \square

Author Contributions K. Ariu, A. Proutiere, and S. Yun collaboratively developed the conceptualization and formulation of the model and problem. The establishment of proofs was carried out by K. Ariu, J. Ok, and S. Yun. Both K. Ariu and J. Ok conducted numerical experiments using synthetic data, while K. Ariu also conducted experiments with non-synthetic data. All authors actively participated in writing the entire manuscript.

Funding Open access funding provided by Royal Institute of Technology. This research is partly funded by the Nakajima Foundation Scholarship (Kaito Ariu), Vetenskaprådet (Alexandre Proutiere), Institute of Information and Communications Technology Planning and Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST) (Seyoung Yun) and No. 2019-0-01906, Artificial Intelligence Graduate School Program (POSTECH) (Jungseul Ok)), and JSPS KAKENHI Grant No. 23K19986 (Kaito Ariu).

Availability of data and materials We have attached the data used and generated raw data at the following URL: <https://bit.ly/3Am5DhX>.

Code Availability We have attached the source code at the following URL: <https://bit.ly/3Am5DhX>.

Declarations

Conflict of interest Not applicable.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abbe, E. (2018). Community detection and stochastic block models. *Foundations and Trends in Communications and Information Theory*, 14(1–2), 1–162.
- Boys in Bristol Photography. (2024). A Canada goose grooming while swimming in a lake. <https://www.pexels.com/photo/a-canada-goose-grooming-while-swimming-in-a-lake-7589597/>. Online accessed 2 Feb. 2024.
- Chris F. (2024). A mallard duck on water. <https://www.pexels.com/photo/a-mallard-duck-on-water-11798057/>. Online; Accessed 2 Feb. 2024.
- Dawid, A. P., & Skene, A. M. (1979). Maximum likelihood estimation of observer error-rates using the em algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1), 20–28.
- Gao, C., Lu, Y., & Zhou, D. (2016). Exact exponent in optimal rates for crowdsourcing. In *Proceedings of the 33rd international conference on machine learning* (pp. 603–611).
- Gao, C., Ma, Z., Zhang, A. Y., Zhou, H. H., et al. (2018). Community detection in degree-corrected block models. *The Annals of Statistics*, 46(5), 2153–2185.
- Gomes, R. G., Welinder, P., Krause, A., & Perona, P. (2011). Crowdclustering. In *Advances in neural information processing systems*, (pp. 558–566).
- Ho, C.-J., Jabbari, S., & Vaughan, J. W. (2013). Adaptive task assignment for crowdsourced classification. In *International conference on machine learning*, (pp. 534–542).

- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301), 13–30.
- Karger, D. R., Oh, S., & Shah, D. (2011). Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems* (pp. 1953–1961).
- Khetan, A., & Oh, S. (2016). Achieving budget-optimality with adaptive schemes in crowdsourcing. In *Advances in Neural Information Processing Systems* (Vol. 29).
- Lai, T. L., & Herbert, R. (1985). Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1), 4–22.
- Long T.-T., Matteo V., Alex R., & Nicholas R. J. (2013). Efficient budget allocation with accuracy guarantees for crowdsourcing classification tasks. In *Proceedings of the 2013 international conference on autonomous agents and multi-agent systems*, (pp. 901–908)
- Ok, J., Oh, S., Shin, J., & Yi, Y. (2016). Optimality of belief propagation for crowdsourced classification. In *Proceedings of the 33rd International Conference on Machine Learning* (pp. 535–544).
- Raykar, V. C., Shipeng, Yu., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., & Moy, L. (2010). Learning from crowds. *Journal of Machine Learning Research*, 11, 1297–1322.
- Vinayak, R. K., & Hassibi, B. (2016). Crowdsourced clustering: Querying edges vs triangles. In *Advances in Neural Information Processing Systems*, (pp. 1316–1324).
- Welinder, P., Branson, S., Perona, P., & Belongie, S. J. (2010). The multidimensional wisdom of crowds. In *Advances in neural information processing systems*, (pp. 2424–2432).
- Yun, S.-Y., & Proutiere, A. (2016). Optimal cluster recovery in the labeled stochastic block model. In *Advances in Neural Information Processing Systems*, (pp. 965–973).
- Zhang, Y., Chen, X., Zhou, D., & Jordan, M. I. (2014). Spectral methods meet EM: a provably optimal algorithm for crowdsourcing. In *Advances in neural information processing systems*, (pp. 1260–1268).
- Zhou, D., Basu, S., Mao, Y., & Platt, J. C. (2012). Learning from the wisdom of crowds by minimax entropy. In *Advances in neural information processing systems*, (pp. 2195–2203).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Kaito Ariu^{1,2}  · Jungseul Ok³ · Alexandre Proutiere² · Seyoung Yun⁴

✉ Kaito Ariu
kaito_ariu@cyberagent.co.jp

Jungseul Ok
jungseul.ok@postech.ac.kr

Alexandre Proutiere
alepro@kth.se

Seyoung Yun
yunseyoung@kaist.ac.kr

¹ CyberAgent, Tokyo, Japan

² KTH, Stockholm, Sweden

³ POSTECH, Pohang-si, South Korea

⁴ KAIST, Daejeon, South Korea