

Optimal Clustering with Bandit Feedback

Junwen Yang

*Institute of Operations Research and Analytics
National University of Singapore
117602, Singapore*

JUNWEN_YANG@U.NUS.EDU

Zixin Zhong

*Thrust of Data Science and Analytics
Hong Kong University of Science and Technology (Guangzhou)
511453, Guangzhou, Guangdong, China*

ZIXINZHONG@HKUST-GZ.EDU.CN

Vincent Y. F. Tan

*Department of Mathematics
Department of Electrical and Computer Engineering
Institute of Operations Research and Analytics
National University of Singapore
119076, Singapore*

VTAN@NUS.EDU.SG

Editor: Chris Wiggins

Abstract

This paper considers the problem of online clustering with bandit feedback. A set of arms (or items) can be partitioned into various groups that are unknown. Within each group, the observations associated to each of the arms follow the same distribution with the same mean vector. At each time step, the agent queries or pulls an arm and obtains an independent observation from the distribution it is associated to. Subsequent pulls depend on previous ones as well as the previously obtained samples. The agent's task is to uncover the underlying partition of the arms with the least number of arm pulls and with a probability of error not exceeding a prescribed constant δ . The problem proposed finds numerous applications from clustering of variants of viruses to online market segmentation. We present an instance-dependent information-theoretic lower bound on the expected sample complexity for this task, and design a computationally efficient and asymptotically optimal algorithm, namely BANDIT ONLINE CLUSTERING (BOC). The algorithm includes a novel stopping rule for adaptive sequential testing that circumvents the need to exactly solve any NP-hard weighted clustering problem as its subroutines. We show through extensive simulations on synthetic and real-world datasets that BOC's performance matches the lower bound asymptotically, and significantly outperforms a non-adaptive baseline algorithm.

Keywords: clustering, K -means, online learning, multi-armed bandits, pure exploration

1. Introduction

Clustering, the task of partitioning a set of items into smaller clusters based on their commonalities, is one of the most fundamental tasks in data analysis and machine learning with a rich and diverse history (Driver and Kroeber, 1932; Cattell, 1943; Ruspini, 1969; Jain et al., 1999; Celebi et al., 2013). It has numerous applications in a wide variety of

areas including business analytics, bioinformatics, pattern recognition, and social sciences. In this era of abundance of medical data, clustering is a powerful tool to uncover the underlying patterns of unknown treatments or diseases when related systematic knowledge is underdeveloped (MacCuish and MacCuish, 2010). In commercial decision making, aiming at increasing customer satisfaction and maximizing potential benefit, marketers utilize clustering strategies to partition the inclusive business market into more narrow market segments with similar characteristics (Chaturvedi et al., 1997). Due to its enormous importance in practical applications, clustering has been studied extensively in the literature from multidisciplinary perspectives. A plethora of algorithms (e.g., K -means and spectral clustering) have been proposed for the task of clustering (see Jain et al. (1999) or Saxena et al. (2017) for comprehensive reviews). In particular, the K -means algorithm by MacQueen (1967) and Lloyd (1982) is arguably the most ubiquitous algorithm due to its simplicity, efficiency, and empirical successes (Jain, 2010).

Although modern data analysis has benefited immensely from the abundance and richness of data, there is an urgent need to develop new techniques that are adapted to the sequential and uncertain nature of data collection. In this paper, we are interested in *online clustering with bandit feedback*, which is an online variant of the classical offline clustering problem. With bandit feedback, the agent only observes a noisy measurement on the selected arm (or item) at each time step. However, the agent can decide which arm to pull adaptively, so as to minimize the expected number of total arm pulls it takes to correctly partition the given arm set with a given (high) probability.

Two Motivating Examples. Our online clustering model captures various contemporary real-world scenarios, in which data contaminated by some degree of measurement noise become available in a sequential and adaptive fashion. We are firstly motivated by medical and public health professionals’ arduous battles against new viruses (e.g., COVID-19). In the face of an unknown type of virus that has different variants, let us assume that there is only one dominant variant in each sub-region of a particular state. When accurate laboratory analysis is not available especially in underdeveloped regions, how can healthcare professionals partition the virus samples into specific dominant variants based on noisy measurements of infectious patients from various sub-regions? This realistic and critical problem can be well modelled by our framework, namely online clustering with bandit feedback, where the healthcare personnel adaptively obtains independent observations of patients from the selected sub-regions and finally partitions the whole state into various groups based on the types of dominant variants. Due to the prohibitive costs in obtaining the measurements, this must be done with as few of them as possible.

Our second motivation is in digital marketing in which customer feedback on certain products are collected in an online manner and always accompanied by random or systematic noise. For market segmentation, one important objective is to reduce the cost of feedback collection while maintaining a high quality of clustering, so that subsequent recommendations are suitably tailored to particular groups of consumers. See Figure 1 for a protocol of online market segmentation that our framework is able to model well.

Main Contributions. Our main contributions are as follows:

- (i) We formulate the online clustering with bandit feedback problem in Section 3. We identify some subtleties of the framework, which may be of independent interest for




Figure 1: Example involving partitioning 6 sub-groups of customers into 3 market segments with bandit feedback. Initially, customers are preliminarily divided into sub-groups, treated as arms, based on some basic characteristics such as age and gender. At each time step, an algorithm chooses a sub-group to query and receives a multidimensional sample (e.g., product ratings) from a single customer within that sub-group. Based on the previously chosen sub-groups and their samples, the algorithm decides which sub-group to query next. Finally, when it is sufficiently confident of producing a partition of the sub-groups into market segments that share similar preferences, the algorithm terminates.

future research. In particular, there are multiple ways that a partition of an arm set can be represented. This is further complicated by the fact that each cluster is also identified by a mean vector. We propose a precise expression for an instance of cluster bandits, and establish two equivalence relations for the partitions and bandit instances, respectively.

- (ii) In Section 4, we derive an instance-dependent (information-theoretic) lower bound on the expected sample complexity for the online clustering problem; this lower bound, however, involves a tricky optimization problem. By exploiting the structure of the problem and leveraging an interesting combinatorial property, we simplify the optimization to a *finite convex* minimax problem, which can be solved efficiently. Further analyses of the lower bound provide fundamental insights and essential tools for the design of our algorithm.
- (iii) In Section 5, we propose and analyze BANDIT ONLINE CLUSTERING (or BOC). We show that it is not only computationally efficient but also asymptotically optimal in the sense that its expected sample complexity attains the lower bound as the error probability tends to zero. This result is somewhat surprising since solving the corresponding *offline* clustering problem exactly is NP-hard (Aloise et al., 2009; Mahajan et al., 2012). En route to overcoming this combinatorial challenge and demonstrating

desirable properties of BOC, we utilize a variant of the classical K -means algorithm in the sampling rule and propose a novel stopping rule for adaptive sequential testing. In particular, the “natural” stopping rule based on the generalized likelihood ratio (GLR) statistic (Chernoff, 1959; Garivier and Kaufmann, 2016) is intractable. Our workaround involves another statistic that exploits our insights on the lower bound. Finally in Section 6, we show via numerical experiments that BOC is indeed asymptotically optimal and significantly outperforms a uniform sampling strategy on both synthetic and real-world benchmark datasets.

2. Literature Review

Clustering and the K -means algorithm. Although we consider an online version of the clustering problem, a certain variant of the classical (offline) K -means algorithm serves as a subroutine of BOC. Given the d -dimensional observations of M items, the K -means algorithm (MacQueen, 1967; Lloyd, 1982) aims at partitioning the items into K disjoint clusters in order to minimize the sum of squared Euclidean distances between each item to the center of its associated cluster. If each item is associated to a certain *weight*, then a reasonable objective is the minimization of the weighted sum of squared Euclidean distances. Due to the prominence of the K -means algorithm, the corresponding clustering problem is also referred to as the (*weighted*) K -means clustering problem in the literature. This non-convex K -means clustering problem has been shown to be NP-hard even for $d = 2$ (Mahajan et al., 2012) or $K = 2$ (Aloise et al., 2009). In fact, for general d , K and M , the problem can be exactly solved in $O(M^{dK+1})$ time (Inaba et al., 1994). When used as a heuristic algorithm, the performance of K -means depends to a large extent on how it is initialized. To improve the stability and the quality of the eventual solution that K -means produces, various initialization methods have been proposed (e.g., Forgy’s method (Forgy, 1965), Maximin (Gonzalez, 1985), K -means++ (Arthur and Vassilvitskii, 2007), PCA-Part (Su and Dy, 2007)). See Celebi et al. (2013) for detailed comparisons on initialization methods for K -means.

There have been some attempts in the literature to adapt the vanilla K -means algorithm to an online framework involving streams of incoming data (Choromanska and Monteleoni, 2012; Liberty et al., 2016; Cohen-Addad et al., 2021). In this line of works, each item only has *one* observation. In particular, at each time step t , the agent receives an observation of one item and has to determine its cluster index before the arrival of next observation. This is vastly different from our setting of bandit feedback, where the arm pulls are determined adaptively by the agent and the observations are stochastic. In addition, another related work by Khaleghi et al. (2012) addressed the case where every item is associated with an infinite sequence generated by one of K unknown stationary ergodic processes. At each time step t , some observations arrive, each being either a new sequence or the continuation of some previously observed sequence, and the agent needs to partition the observed sequences into K groups. Finally, Mazumdar and Saha (2017) studied clustering with noisy queries in both the adaptive and non-adaptive settings, where the agent receives a potentially incorrect answer on whether two items belong to the same cluster at each time step. To the best of our knowledge, the online clustering with bandit feedback problem (formally described in Section 3) has not been considered before.

Bandit Algorithms. The stochastic multi-armed bandit problem, originally introduced by Thompson (1933), provides a simple but powerful online learning framework. This problem has been studied extensively in recent years. While the *regret minimization* problem aims at maximizing the cumulative reward by balancing the trade-off between exploration and exploitation (Auer et al., 2002; Abbasi-Yadkori et al., 2011; Bubeck and Cesa-Bianchi, 2012; Agrawal and Goyal, 2012), the *pure exploration* problem focuses on efficient exploration with specific goals, e.g., best (top- k) arm identification (Even-Dar et al., 2006; Audibert et al., 2010; Karnin et al., 2013; Garivier and Kaufmann, 2016; Jun et al., 2016), and its variants in linear bandits (Soare et al., 2014; Jedra and Proutiere, 2020; Yang and Tan, 2022), and cascading bandits (Zhong et al., 2020), among others. Our online clustering task can also be viewed as a pure exploration problem although the rewards are *multi-dimensional* and the arm set has an inherent cluster structure. It is worth mentioning that Prabhu et al. (2020) introduced a framework of sequential multi-hypothesis testing with bandit feedback, which is a generalization of the odd arm identification problem (Vaidhiyan and Sundaresan, 2017). Our online clustering problem falls within this framework if each partition of the arm set is viewed as a hypothesis. However, the methodology proposed in Prabhu et al. (2020) relies heavily on a strong continuity assumption on the proportions of arm pulls (Assumption A therein). Even if one accepts continuity assumption, the total number of hypotheses (i.e., the number of possible partitions) is prohibitively large and the corresponding computation of the *modified GLR* is intractable (see Remark 17 for more details). Our work does away with the continuity assumption and, in fact, *proves* that the required continuity property holds (see Proposition 9). We refer to Lattimore and Szepesvári (2020) for a comprehensive review on bandit algorithms.

There is also a line of works that incorporates cluster structures into multi-armed bandits. In particular, Nguyen and Lauw (2014); Gentile et al. (2014), Li et al. (2016), and Carlsson et al. (2021) assumed that users can be divided into groups and the users within each group receive similar rewards for each arm. Besides, Bouneffouf et al. (2019) and Singh et al. (2020) assumed that the arm set is pre-clustered and the reward distributions of the arms within each cluster are similar. However, all the works mentioned above focus on *leveraging* the cluster structure to improve the performance of regret minimization, which differs from our objective, i.e., to *uncover* the underlying partition of the arms. Finally, Wang and Scarlett (2022) considered a min-max grouped bandits problem in which the objective is to find a sub-group (among possibly overlapping groups) whose worst arm has the highest mean reward.

3. Problem Setup and Preliminaries

A Bandit Feedback Model with Cluster Structure. We consider a bandit feedback model, in which the arm set has an inherent cluster structure. In particular, the agent is given an arm set $\mathcal{A} = [M]$, which can be partitioned into K disjoint nonempty clusters, and the arms in the same cluster share the same d -dimensional mean vector, also referred to as the *center* of the corresponding cluster. Without loss of generality, we assume that $K < M$, otherwise there is only one possible partition. Therefore, an instance of cluster bandits can be fully characterized by a pair (c, \mathcal{U}) , where $c = [c_1, c_2, \dots, c_M] \in [K]^M$ consists of the cluster indices of the arms and $\mathcal{U} = [\mu(1), \mu(2), \dots, \mu(K)] \in \mathbb{R}^{d \times K}$ represents the K centers




Figure 2: Online clustering with bandit feedback with $K = 3$ and $M = 12$.

of the clusters. Since each cluster has at least one arm, for any cluster $k \in [K]$, there exists an arm $m \in [M]$ such that $c_m = k$. To reduce clutter and ease the reading, we always index the arms and the clusters by subscripts and numbers in parentheses, respectively.

At each time t , the agent selects an arm A_t from the arm set \mathcal{A} , and then observes a noisy measurement on the mean vector of A_t , i.e.,

$$X_t = \mu(c_{A_t}) + \eta_t$$

where $\{\eta_t\}_{t=1}^{\infty} \subset \mathbb{R}^d$ is a sequence of independent (noise) random variables, each following the standard d -dimensional Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. See Figure 2 for a schematic of the model.

Remark 1 *In this work, we assume that the noise at each time t is a standard d -dimensional Gaussian random vector. For general (e.g., non-diagonal) noise covariance matrices, one can use the Cholesky decomposition to transform the raw observations in an affine manner into ones that have an identity covariance matrix, without loss of generality.*

The Equivalences of Partitions and Instances. The representation of a partition c or an instance (c, \mathcal{U}) is not unique and we can accordingly define two equivalence relations. For a permutation σ on $[K]$, let $\sigma(c) := [\sigma(c_1), \sigma(c_2), \dots, \sigma(c_M)]$ and $\sigma(\mathcal{U}) := [\mu(\sigma(1)), \mu(\sigma(2)), \dots, \mu(\sigma(K))]$. Similarly, we define $\sigma(c, \mathcal{U}) := (\sigma(c), \sigma(\mathcal{U}))$.

For two partitions c and c' , if there exists a permutation σ on $[K]$ such that $c = \sigma(c')$, then we write $c \sim c'$. Due to the bijectivity of permutations, there also exists another permutation σ^{-1} on $[K]$ such that $c' = \sigma^{-1}(c)$. Therefore, it is straightforward to verify this is indeed an equivalence relation.

For two instances (c, \mathcal{U}) and (c', \mathcal{U}') ,¹ if $\mu(c_m) = \mu'(c'_m)$ for all $m \in [M]$, then the two instances are equivalent and we denote this as $(c, \mathcal{U}) \sim (c', \mathcal{U}')$. Note that $(c', \mathcal{U}') = \sigma(c, \mathcal{U})$ for some permutation σ indicates $(c, \mathcal{U}) \sim (c', \mathcal{U}')$, but the reverse implication may not be true. That is to say, if $(c, \mathcal{U}) \sim (c', \mathcal{U}')$, there may not exist a permutation σ such that $(c', \mathcal{U}') = \sigma(c, \mathcal{U})$. See Example 1.

1. Throughout this paper, we denote the columns of $\mathcal{U}' \in \mathbb{R}^{d \times K}$ as $\{\mu(k)\}_{k=1}^K$. The same convention applies to other forms of \mathcal{U} , such as $\bar{\mathcal{U}}$, \mathcal{U}^* , and \mathcal{U}^\dagger .

Example 1 Let $K = 2$, $M = 3$ and $d = 1$. Consider the three instances (c, \mathcal{U}) , (c', \mathcal{U}') and (c'', \mathcal{U}'') where $c = [1, 1, 2]$, $c' = [2, 2, 1]$, $c'' = [1, 2, 2]$ and $\mathcal{U} = \mathcal{U}' = \mathcal{U}'' = [1, 1]$. Although $(c, \mathcal{U}) \sim (c', \mathcal{U}') \sim (c'', \mathcal{U}'')$, it holds that $c \sim c' \not\sim c''$.

Online Clustering with Bandit Feedback. We consider a pure exploration task, aiming to find the unknown cluster structure c by pulling arms adaptively. In the fixed-confidence setting where a confidence level $\delta \in (0, 1)$ is given, the agent is required to find a *correct* partition c^{out} of the arm set $[M]$ (i.e., $c^{\text{out}} \sim c$) with a probability of at least $1 - \delta$ in the smallest number of time steps.

More formally, the agent uses an *online algorithm* π to decide the arm A_t to pull at each time step t , to choose a time τ_δ to stop pulling arms, and to recommend c^{out} as the partition to output eventually. Let $\mathcal{F}_t = \sigma(A_1, X_1, \dots, A_t, X_t)$ denote the σ -field generated by the past measurements up to and including time t . Thus, the online algorithm π consists of three components, namely,

- the *sampling rule* selects A_t , which is \mathcal{F}_{t-1} -measurable;
- the *stopping rule* determines a stopping time τ_δ adapted to the filtration $(\mathcal{F}_t)_{t=1}^\infty$;
- the *recommendation rule* outputs a partition c^{out} , which is $\mathcal{F}_{\tau_\delta}$ -measurable.

Definition 2 For a fixed confidence level $\delta \in (0, 1)$, an online clustering algorithm π is said to be δ -PAC (probably approximately correct) if for all instances (c, \mathcal{U}) , $\Pr(\tau_\delta < \infty) = 1$ and the probability of error $\Pr(c^{\text{out}} \not\sim c) \leq \delta$.

Our overarching goal is to design a computationally efficient online δ -PAC clustering algorithm while minimizing the expected sample complexity $\mathbb{E}[\tau_\delta]$. To rule out pathological cases that might lead to infinite expected sample complexities for any algorithm, throughout this work, we only consider partitioning the instances (c, \mathcal{U}) that satisfy the following natural property: the mean vectors for different clusters are *distinct* (i.e., the instances (c, \mathcal{U}) subject to $\mathcal{U} \in \mathcal{U} := \{\bar{\mathcal{U}} \in \mathbb{R}^{d \times K} : \bar{\mu}(k_1) \neq \bar{\mu}(k_2) \text{ for all } 1 \leq k_1 < k_2 \leq K\}$). However, this property might not hold for general instances of cluster bandits (including the alternative instances $\text{Alt}(c)$ that we introduce in the next section). Note that the odd arm identification problem (Vaidhiyan and Sundaresan, 2017; Karthik and Sundaresan, 2020, 2021) is not a special case of our online clustering problem with $K = 2$ since we do not require the knowledge of the number of arms in each cluster.

Other Notations. Let \mathbb{N} denote the set of positive integers and \mathbb{R}^+ denote the set of non-negative real numbers. For any positive integer N , $\mathcal{P}_N := \{x \in [0, 1]^N : \|x\|_1 = 1\}$ denotes the probability simplex in \mathbb{R}^N while $\mathcal{P}_N^+ := \{x \in (0, 1)^N : \|x\|_1 = 1\}$ denotes the open probability simplex in \mathbb{R}^N .

For two partitions c and c' , let $d_H(c, c')$ denote the Hamming distance between c and c' , i.e., $d_H(c, c') := \sum_{m=1}^M \mathbb{1}\{c_m \neq c'_m\}$. For any $a, b \in (0, 1)$, the binary relative entropy, which is the KL-divergence between Bernoulli distributions with means a and b , is denoted as $d_{\text{KL}}(a, b) := a \log(a/b) + (1 - a) \log((1 - a)/(1 - b))$.

When we write $i^* = \arg \min_{i \in A} f(i)$ where $f(i)$ is a function of i and A is a finite set of integers or a finite set of vectors of integers, we are referring to the minimum index (in lexicographic order) in the set $\{i \in A : f(i) = \min_{j \in A} f(j)\}$ if it is not a singleton.

4. Lower Bound

In this section, we leverage the ubiquitous change-of-measure argument for deriving impossibility results to derive an instance-dependent lower bound on the expected sample complexity $\mathbb{E}[\tau_\delta]$ for the online clustering problem. The lower bound is closely related to a combinatorial optimization problem. Although the optimization in its original form appears to be intractable, we prove an interesting combinatorial property and reformulate the optimization as a finite convex minimax problem. Moreover, we further present some results on the computation and other useful properties (e.g., the continuity of the optimizer and the optimal value) of the optimization problem (and its sub-problem) embedded in the lower bound, which are fundamental and essential in our algorithm design (see Section 5).

The change-of-measure argument, of which the key idea dates back to Chernoff (1959), is ubiquitous in showing various lower bounds in (and beyond) bandit problems (e.g., regret minimization (Lai and Robbins, 1985; Lattimore and Szepesvári, 2017), pure exploration (Kaufmann et al., 2016; Vaidhiyan and Sundaresan, 2017)). Using this technique, the probabilities of the same event under different probability measures are related via the KL-divergence between the two measures.

For any fixed instance (c, \mathcal{U}) , we define $\text{Alt}(c) := \{(c', \mathcal{U}') : c' \neq c \text{ for any } (c'', \mathcal{U}'') \sim (c', \mathcal{U}')\}$, which is the set of alternative instances where c is not a correct partition. In particular, we consider the probabilities of correctly identifying the cluster structures under (c, \mathcal{U}) and the instances in $\text{Alt}(c)$, and apply the *transportation* inequality (Kaufmann et al., 2016, Lemma 1). The instance-dependent lower bound on $\mathbb{E}[\tau_\delta]$ is presented in the following theorem; see Appendix B.1 for the proof.

Theorem 3 *For a fixed confidence level $\delta \in (0, 1)$ and instance (c, \mathcal{U}) , any δ -PAC online clustering algorithm satisfies*

$$\mathbb{E}[\tau_\delta] \geq d_{\text{KL}}(\delta, 1 - \delta) D^*(c, \mathcal{U})$$

where

$$D^*(c, \mathcal{U}) := \left(\frac{1}{2} \sup_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \right)^{-1}. \quad (1)$$

Furthermore,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)} \geq D^*(c, \mathcal{U}). \quad (2)$$

We refer to $D^*(c, \mathcal{U})$ as the *hardness parameter* of the online clustering task in the sequel. The asymptotic version of the instance-dependent lower bound given in Equation (2) in Theorem 3 is tight in view of the expected sample complexity of the efficient algorithm we present in Section 5. Intuitively, any $\lambda \in \mathcal{P}_M$ in Equation (1) can be understood as the proportion of arm pulls, which inspires the design of the sampling rule of a δ -PAC online clustering algorithm. The agent wishes to find the optimal proportion of arm pulls to distinguish the instance c from the *most confusing* alternative instances in $\text{Alt}(c)$ (for which c is not a correct partition). Therefore, with the knowledge of the instance (c, \mathcal{U}) , the

optimization problem embedded in (1) naturally unveils the optimal sampling rule, which is the basic idea behind the design of our sampling rule in Section 5.

For ease of description, we refer to the entire optimization problem and the inner infimization in (1) as Problem (\square) and Problem (\triangle) , respectively, i.e.,

$$\begin{aligned} \text{Problem } (\square): \quad & \sup_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \\ \text{Problem } (\triangle): \quad & \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2. \end{aligned}$$

These two optimization problems will also feature in the sampling rule and the stopping rule, respectively, of our proposed algorithm in Section 5. However, even given the full information of the instance (c, \mathcal{U}) , solving Problem (\square) is tricky although the inner objective function is a weighted sum of quadratic functions. In particular, the alternative instances in $\text{Alt}(c)$ need to be identified by *both* c and \mathcal{U}' , which is rather involved as the definition of $\text{Alt}(c)$ is combinatorial and the number of instances in it is obviously infinite. For Problem (\triangle) , one may consider first fixing c' , then optimizing over different \mathcal{U}' . We remark that this idea is only theoretically but not practically feasible since for a fixed number of clusters K , the total number of possible partitions (which is called *the Stirling number of the second kind* (Graham et al., 1989)) grows asymptotically as $K^M/K!$ as the number of arms M tends to infinity. Nevertheless, we show a natural and useful property of Problem (\triangle) .

Lemma 4 *For any $\lambda \in \mathcal{P}_M$ and (c, \mathcal{U}) ,*

$$\inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 = \inf_{\substack{(c', \mathcal{U}') \in \text{Alt}(c): \\ d_{\text{H}}(c', c) = 1}} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2.$$

Lemma 4 provides a useful combinatorial property of Problem (\triangle) , which makes it possible to solve for the hardness parameter $D^*(c, \mathcal{U})$ efficiently. Instead of considering *all* the alternative instances in $\text{Alt}(c)$, Lemma 4 shows it suffices to consider the instances whose partitions have a *Hamming distance of 1* from the given partition c . The complete proof of Lemma 4 is deferred to Appendix B.2, which consists of four steps (see Figure 4 therein for an illustration). In particular, we show for any instance $(c^\dagger, \mathcal{U}^\dagger) \in \text{Alt}(c)$ such that $d_{\text{H}}(c^\dagger, c) > 1$, there exists another instance $(c^*, \mathcal{U}^*) \in \text{Alt}(c)$ such that $d_{\text{H}}(c^*, c) = 1$ and the objective function under (c^*, \mathcal{U}^*) is not larger than that under $(c^\dagger, \mathcal{U}^\dagger)$. This desirable combinatorial property depends strongly on the specific structure of a valid partition. In fact, in Example 2 in Section 5, we will see an example where a similar combinatorial property no longer holds if the true mean vectors $\{\mu(c_m)\}_{m=1}^M$ in (1) are replaced by some empirical estimates which may be obtained as the algorithm proceeds. Thanks to Lemma 4, Problem (\triangle) turns out to be equivalent to a much simpler finite minimization problem, as shown in Proposition 5.

Proposition 5 For any $\lambda \in \mathcal{P}_M$ and (c, \mathcal{U}) ,

$$\inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 = \begin{cases} \min_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \frac{\bar{w}(k)w(k')}{\bar{w}(k)+w(k')} \|\mu(k) - \mu(k')\|^2 & \text{if } \lambda \in \mathcal{P}_M^+ \\ 0 & \text{otherwise} \end{cases}$$

where $w(k) := \sum_{m=1}^M \lambda_m \mathbb{1}\{c_m = k\}$, $n(k) := \sum_{m=1}^M \mathbb{1}\{c_m = k\}$ and $\bar{w}(k) := \min_{m \in [M]: c_m = k} \lambda_m$. Moreover, if $\lambda \in \mathcal{P}_M^+$, the infimum in Problem (Δ) can be replaced with a minimum.

As a corollary of our former intuition where λ represents the proportion of arm pulls, given the knowledge of the instance (c, \mathcal{U}) and the proportion of arm pulls, Problem (Δ) tells us how similar the true instance c and the most confusing alternative instances in $\text{Alt}(c)$ are. In fact, Proposition 5 plays an essential role in the computation of the stopping rule of our method in Section 5, which succeeds in circumventing the need to solve NP-hard optimization problems. In addition, Proposition 6 below asserts the continuity of the optimal value of Problem (Δ) , which will help to assert that the stopping rule proposed in Section 5 is asymptotically optimal. Refer to Appendices B.3 and B.4 for the proofs of Proposition 5 and Proposition 6, respectively.

Proposition 6 For any fixed c , define $g : \mathcal{P}_M \times \mathbb{R}^{d \times K} \rightarrow \mathbb{R}^+$ as

$$g(\lambda, \mathcal{U}) := \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2.$$

Then g is continuous on $\mathcal{P}_M \times \mathcal{U}$.

As a consequence of Proposition 5, Proposition 7 below (proved in Appendix B.5), transforms Problem (\square) into a finite convex minimax problem, which has been studied extensively in the optimization literature (e.g., Gigola and Gomez (1990) Herrmann (1999), Gaudioso et al. (2006)).

Proposition 7 For any (c, \mathcal{U}) ,

$$D^*(c, \mathcal{U}) = 2 \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2}.$$

A by-product of Proposition 7 is that the outer supremum in Problem (\square) can be replaced with a maximum. Intuitively, the maximizer of Problem (\square) in \mathcal{P}_M represents the optimal proportion of arm pulls, which will be of considerable importance in our design of the sampling rule. In fact, there exists a bijection between the solution to the finite convex minimax problem above and Problem (\square) , as shown in Proposition 8 below. Although the finite convex minimax problem is not strictly convex in $w \in \mathcal{P}_K^+$, Proposition 8 states that the solution to Problem (\square) is unique. This, together with Proposition 9 concerning the continuity of the solution to Problem (\square) , guarantees the computational efficiency and the asymptotic optimality of our sampling rule in Section 5. The proofs of Proposition 8 and Proposition 9 are deferred to Appendices B.6 and B.7, respectively.

Proposition 8 For any (c, \mathcal{U}) , the solution to

$$\arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \quad (3)$$

is unique.

If λ^* denotes the unique solution to (3) and w^* denotes the unique solution to

$$\arg \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2},$$

then λ^* can be expressed in terms of w^* as

$$\lambda_m^* = \frac{w^*(c_m)}{n(c_m)} \text{ for all } m \in [M].$$

Proposition 9 For any fixed c , define $\Lambda : \mathbb{R}^{d \times K} \rightarrow \mathcal{P}_M$ as

$$\Lambda(\mathcal{U}) := \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2.$$

Then Λ is continuous² on \mathcal{U} .

5. Algorithm: Bandit Online Clustering

For the online clustering task with bandit feedback, we propose a computationally efficient and asymptotically optimal algorithm, namely BANDIT ONLINE CLUSTERING (or BOC), whose pseudocode is presented in Algorithm 1. As elucidated in Section 3, our algorithm comprises three components: the sampling rule for selecting the arms to pull (Lines 3-9), the stopping rule for determining a stopping time (Lines 11-13), and the recommendation rule for producing a final result (the clustering c^{out}). In the following subsections, we first introduce an important subroutine called K-MEANS-MAXIMIN. This subroutine is employed at each time step to return a guess of both the correct partition and mean vectors based on the past measurements. We then explain the methodology of our sampling rule and stopping rule, respectively, focusing on their connections with the optimization problems discussed in Section 4. Finally, from a theoretical standpoint, we demonstrate the effectiveness of BOC and prove that its expected sample complexity $\mathbb{E}[\tau_\delta]$ is asymptotically optimal as the confidence level δ tends to zero.

5.1 Weighted K -means with Maximin Initialization

Although we only aim at producing a correct partition in the final recommendation rule as noted in Section 3, learning the K unknown mean vectors of the clusters is essential in the sampling rule as well as the stopping rule. Different from other pure exploration problems in bandits (e.g., best arm identification (Garivier and Kaufmann, 2016), odd

2. In finite-dimensional spaces, pointwise convergence and convergence in L^p norm are equivalent.

Algorithm 1 Bandit Online Clustering (BOC)

Input: Number of clusters K , confidence level δ and arm set $[M]$

 1: Sample each arm once, set $t = M$ and initialize $\hat{\mu}_m(t)$ and $N_m(t) = 1$ for all $m \in [M]$.

 2: **repeat**

 3: **if** $\min_{m \in [M]} N_m(t) \leq \max(\sqrt{t} - M/2, 0)$ **then**

 4: Sample $A_{t+1} = \arg \min_{m \in [M]} N_m(t)$, and $(c^t, \mathcal{U}^t) \leftarrow (c^{t-1}, \mathcal{U}^{t-1})$
▷ Forced exploration

 5: **else**

 6: $(c^t, \mathcal{U}^t) \leftarrow \text{K-MEANS-MAXIMIN}(K, \{\hat{\mu}_m(t)\}_{m \in [M]}, \{N_m(t)\}_{m \in [M]})$ ▷ Algorithm 2

 7: Solve ▷ Proposition 8

$$\lambda^*(t) = \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c^t)} \sum_{m=1}^M \lambda_m \|\mu^t(c_m^t) - \mu'(c'_m)\|^2 \quad (4)$$

 8: Sample $A_{t+1} = \arg \max_{m \in [M]} (t\lambda_m^*(t) - N_m(t))$

 9: **end if**

 10: $t \leftarrow t + 1$, update $\hat{\mu}_m(t)$ and $N_m(t)$ for all $m \in [M]$

11: Compute

$$Z_1(t) = \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2$$

and solve

▷ Proposition 5

$$Z_2(t) = \min_{(c', \mathcal{U}') \in \text{Alt}(c^{t-1})} \sum_{m=1}^M N_m(t) \|\mu^{t-1}(c_m^{t-1}) - \mu'(c'_m)\|^2$$

12: Set

$$Z(t) = \frac{1}{2} \left(\left(-\sqrt{Z_1(t)} + \sqrt{Z_2(t)} \right)_+ \right)^2$$

 13: **until** $Z(t) \geq \beta(\delta, t)$
Output: $\tau_\delta = t$ and $c^{\text{out}} = c^{t-1}$

arm identification (Vaidhiyan and Sundaresan, 2017)), in the online clustering problem, it is not straightforward to recommend an estimate of the pair (c, \mathcal{U}) given some past measurements on the arm set. However, using maximum likelihood estimation, we will see the equivalence between the recommendation subroutine and the classical offline *weighted K-means clustering* problem, which has been shown to be NP-hard (Aloise et al., 2009; Mahajan et al., 2012).

Given the past arm pulls and observations up to time t (i.e., $A_1, X_1, \dots, A_t, X_t$), the log-likelihood function of the hypothesis that the instance can be identified by the pair (c', \mathcal{U}') can be written as

$$\ell(c', \mathcal{U}' \mid A_1, X_1, \dots, A_t, X_t) := -\frac{1}{2} \sum_{s=1}^t \|X_s - \mu'(c'_{A_s})\|^2 - \frac{td}{2} \log(2\pi). \quad (5)$$

For any arm $m \in [M]$, let $N_m(t) := \sum_{s=1}^t \mathbb{1}\{A_s = m\}$ and $\hat{\mu}_m(t) := \sum_{s=1}^t X_s \mathbb{1}\{A_s = m\} / N_m(t)$ denote the number of pulls and the empirical estimate up to time t , respectively. By rearranging Equation (5), the maximum likelihood estimate of the unknown pair (c, \mathcal{U}) can be expressed as

$$\arg \min_{(c', \mathcal{U}')} \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu'(c'_m)\|^2 \quad (6)$$

which consists in minimizing a weighted sum of squared Euclidean distances between the empirical estimate of each arm and its associated center. Therefore, any algorithm designed for the weighted K -means clustering problem is applicable to obtain an approximate (not exact) solution to (6).

We remark that although the weighted variant of the original K -means algorithm (MacQueen, 1967; Lloyd, 1982) is an efficient heuristic for the weighted K -means clustering problem, to the best of our knowledge, there are no theoretical guarantees for finding a global minimum of this problem in general. To establish the asymptotic optimality of our online clustering method BOC, in the initialization stage of K -means, we leverage the *Maximin* method, which is a farthest point heuristic proposed by Gonzalez (1985). The complete pseudocode for WEIGHTED K -MEANS WITH MAXIMIN INITIALIZATION (abbreviated as K-MEANS-MAXIMIN) is presented in Algorithm 2. In the following, we derive some useful properties of K-MEANS-MAXIMIN; see Appendix D.1 for the proof of Proposition 10.

Proposition 10 *Given an instance (c, \mathcal{U}) , if the empirical estimates of the arms $\{\hat{\mu}_m\}_{m \in [M]}$ satisfy*

$$\max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| < \frac{1}{4} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|,$$

then K-MEANS-MAXIMIN will output a correct partition $c^{\text{out}} \sim c$. Furthermore, suppose that $c = \sigma(c^{\text{out}})$ for some permutation σ on $[K]$. Then

$$\max_{k \in [K]} \|\mu^{\text{out}}(k) - \mu(\sigma(k))\| \leq \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|.$$

Remark 11 *The K-MEANS-MAXIMIN subroutine used in Algorithm 1 can be replaced by any offline clustering algorithm that meets the following conditions without affecting the subsequent results including Proposition 12, Proposition 18 and Theorem 19: (i) given an instance (c, \mathcal{U}) , if the empirical estimates of all the arms are sufficiently accurate, i.e., $\max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|$ is smaller than a constant ϵ that depends on the problem instance, then the clustering algorithm will output a correct partition c^{out} , i.e., one that satisfies that $c = \sigma(c^{\text{out}})$ for some permutation σ ; (ii) moreover, for any cluster $k \in [K]$, $\|\mu^{\text{out}}(k) - \mu(\sigma(k))\|$ is not larger than $\max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|$.*

5.2 Sampling Rule

Once the estimates of the partition and the mean vectors are obtained by the K-MEANS-MAXIMIN subroutine based on the past measurements, our sampling rule utilizes the efficient method presented in Proposition 8 to find a plug-in approximation of the optimal oracle sampling rule. This then informs the algorithm of the selection of the next arm to pull.

Algorithm 2 Weighted K -means with Maximin Initialization (K-MEANS-MAXIMIN)

Input: Number of clusters K , empirical estimate $\hat{\mu}_m$ and weighting N_m for all $m \in [M]$

- 1: Choose the empirical estimate of an arbitrary arm as the first cluster center $\hat{\mu}(1)$
- 2: **for** $k = 2$ **to** K **do** ▷ Maximin Initialization
- 3: Choose the empirical estimate of the arm that has the greatest Euclidean distance to the nearest existing center as the k -th center $\hat{\mu}(k)$:

$$\hat{\mu}(k) = \arg \max_{m \in [M]} \min_{1 \leq k' \leq k-1} \|\hat{\mu}_m - \hat{\mu}(k')\|$$

- 4: **end for**
- 5: **repeat** ▷ Weighted K -means
- 6: Assign each arm to its closest cluster center:

$$\hat{c}_m = \arg \min_{k \in [K]} \|\hat{\mu}_m - \hat{\mu}(k)\|$$

- 7: Update each cluster center as the weighted mean of the empirical estimates of the arms in it:

$$\hat{\mu}(k) = \frac{\sum_{m \in [M]} N_m \hat{\mu}_m \mathbb{1}\{\hat{c}_m = k\}}{\sum_{m \in [M]} N_m \mathbb{1}\{\hat{c}_m = k\}}$$

- 8: **until** Clustering \hat{c} no longer changes
 - 9: Set $\mu^{\text{out}}(k) = \hat{\mu}(k)$ for all $k \in [K]$
- Output:** $c^{\text{out}} = \hat{c}$ and $\mathcal{U}^{\text{out}} = [\mu^{\text{out}}(1), \mu^{\text{out}}(2), \dots, \mu^{\text{out}}(K)]$
-

In particular, Algorithm 1 follows the so-called *D-Tracking* rule, originally proposed by Garivier and Kaufmann (2016) for best arm identification, to track the optimal sampling rule. For the purpose of ensuring the plug-in approximation $\lambda^*(t)$ to converge to the true optimal sampling rule λ^* , the D-Tracking rule introduce a stage of *forced exploration*. At each time t , if there exists an arm whose number of pulls is not larger than the preset threshold ($\max(\sqrt{t} - M/2, 0)$), then the agent chooses to sample that under-sampled arm. Otherwise, the agent chooses the arm according to the difference between the current proportions of arm pulls and the plug-in approximation $\lambda^*(t)$ (as described in Line 8 of Algorithm 1). Proposition 12, proved in Appendix D.2, shows the asymptotic optimality of our sampling rule, which is a joint consequence of Propositions 8, 9 and 10.

Proposition 12 *Using the sampling rule as detailed in Lines 3-9 of Algorithm 1, the proportions of arm pulls in Algorithm 1 converge to the optimal oracle sampling rule λ^* almost surely, i.e.,*

$$\Pr \left(\lim_{t \rightarrow \infty} \frac{N_m(t)}{t} = \lambda_m^* \text{ for all } m \in [M] \right) = 1.$$

Remark 13 *We utilize the forced exploration stage in Algorithm 1 (Lines 3 and 4) to obtain coarse estimates of the means of the arms. This appears to be necessary since we do not make any assumption on the instance of cluster bandits except that the centers for different clusters are distinct. However, the clusters being distinct does not preclude them being arbitrarily*

close to one another. This hinders the adoption of a “more adaptive” exploration stage. When the ratio between the minimal and the maximal pairwise distances among the centers is lower bounded by a known constant $0 < r \leq 1$ (see Assumption 1), an improved and more aggressive exploration method is presented in Equation (22) of Appendix C. This is based on a delicate quantitative analysis of λ^* (Proposition 28 therein). Experimental results using this more aggressive forced exploration procedure are also presented in Appendix C.2.

5.3 Stopping Rule

As the arm sampling proceeds, the algorithm needs to determine when to stop the sampling and recommend a partition with an error probability of at most δ , namely the stopping rule. Most existing algorithms for pure exploration in the fixed-confidence setting (e.g., Garivier and Kaufmann (2016), Jedra and Proutiere (2020), Feng et al. (2021), Réda et al. (2021)) consider the *Generalized Likelihood Ratio* (GLR) statistic and find suitable task-specific threshold functions. This strategy dates back to Chernoff (1959). However, we show that the method based on the standard GLR is computationally intractable in the following.

Let $(c^{t*}, \mathcal{U}^{t*})$ be the maximum likelihood estimate of the unknown pair (c, \mathcal{U}) given the past measurements up to time t . Then $(c^{t*}, \mathcal{U}^{t*})$ is also the global minimizer to (6), i.e., $(c^{t*}, \mathcal{U}^{t*}) = \arg \min_{(c', \mathcal{U}')} \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu'(c'_m)\|^2$. Using the definition of the log-likelihood function in (5), the logarithm of the GLR statistic (referred to as the log-GLR) for testing $(c^{t*}, \mathcal{U}^{t*})$ against its alternative instances can be written as

$$\begin{aligned} \log\text{-GLR} &= \ell(c^{t*}, \mathcal{U}^{t*} \mid A_1, X_1, \dots, A_t, X_t) - \min_{(c', \mathcal{U}') \in \text{Alt}(c^{t*})} \ell(c', \mathcal{U}' \mid A_1, X_1, \dots, A_t, X_t) \\ &= \frac{1}{2} \left(- \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu^{t*}(c_m^{t*})\|^2 + \min_{(c', \mathcal{U}') \in \text{Alt}(c^{t*})} \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu'(c'_m)\|^2 \right). \end{aligned} \quad (7)$$

There are two critical computational issues with the above expression. First, evaluating the log-GLR as Equation (7) requires the exact global minimizer to (6), which is NP-hard to find in general. More importantly, even with the knowledge of $(c^{t*}, \mathcal{U}^{t*})$, one cannot efficiently solve the optimization problem in the second term of Equation (7) although it appears to be similar to Problem (Δ) discussed in Section 4. One may conjecture that an analogous combinatorial property to Lemma 4 holds for the second term of Equation (7), which might shrink the feasible set from $\text{Alt}(c^{t*})$ to the alternative instances whose partitions have a Hamming distance of exactly 1 from c^{t*} . We disprove this conjecture by showing a counterexample in Example 2. As a consequence, the only feasible approach, at least for the moment, is to check all the possible partitions in $\text{Alt}(c^{t*})$, which is certainly computationally intractable.

Example 2 Let $K = 2$, $M = 4$ and $d = 2$. At time t , suppose that the empirical estimates of the 4 arms are respectively $[0, 0]^\top$, $[a, 0]^\top$, $[0, b]^\top$ and $[a, b]^\top$, where $0 < b/\sqrt{2} < a < b$. Then the partition that attains the minimum in (6) is $[1, 1, 2, 2]$ (or $[2, 2, 1, 1]$) whereas the partition of the most confusing alternative instances (the minimizer to the second term of Equation (7)) is $[1, 2, 1, 2]$ (or $[2, 1, 2, 1]$). Obviously, their Hamming distance is 2 rather than 1.

To construct a practicable stopping rule, instead of the GLR statistic, we consider the statistic

$$Z(t) := \frac{1}{2} \left(\left(-\sqrt{Z_1(t)} + \sqrt{Z_2(t)} \right)_+ \right)^2$$

with

$$Z_1(t) := \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2$$

and

$$Z_2(t) := \min_{(c', \mathcal{U}') \in \text{Alt}(c^{t-1})} \sum_{m=1}^M N_m(t) \|\mu^{t-1}(c_m^{t-1}) - \mu'(c'_m)\|^2.$$

Note that $Z(t)$ involves $(c^{t-1}, \mathcal{U}^{t-1})$, which is the estimate of the true pair (c, \mathcal{U}) produced by the K-MEANS-MAXIMIN subroutine based on the past measurements. In particular, $(c^{t-1}, \mathcal{U}^{t-1})$ is not necessarily the global minimizer to (6) and our stopping rule does not have any requirement on the quality of this estimate. For the term $Z_2(t)$, Proposition 5 can be utilized directly since the inherent optimization is equivalent to Problem (Δ) after an appropriate normalization.

Let $\zeta(\cdot)$ denote the Riemann zeta function, i.e., $\zeta(s) = \sum_{n=1}^{\infty} n^{-s}$. The stopping time of Algorithm 1 is defined as

$$\tau_\delta := \inf\{t \in \mathbb{N} : Z(t) \geq \beta(\delta, t)\}$$

where

$$\beta(\delta, t) = \sum_{m=1}^M 2d \log(4 + \log(N_m(t))) + Md \cdot \Psi \left(\frac{\log(1/\delta)}{Md} \right) \quad (8)$$

with

$$\Psi(x) = \min_{1/2 \leq h \leq 1} \left(2 - 2 \log(4h) + \frac{\log(\zeta(2h))}{h} - \frac{\log(1-h)}{2h} + \frac{x}{h} \right)$$

is a threshold function inspired by the concentration results for univariate Gaussian distributions (Kaufmann and Koolen, 2021).

Remark 14 *The function Ψ used in the threshold $\beta(\delta, t)$ possesses some useful properties that can be verified in a straightforward manner or found in Kaufmann and Koolen (2021): (i) $\Psi(x) = x + \log(x) + o(\log(x))$ as $x \rightarrow \infty$; (ii) $\Psi(x) \geq x$ for all $x > 0$; (iii) let $\psi(h) := 2 - 2 \log(4h) + \frac{\log(\zeta(2h))}{h} - \frac{\log(1-h)}{2h} + \frac{x}{h}$ and then $-\psi$ is a unimodal function on $[1/2, 1]$.*

Overall, our stopping rule is easy to implement and computationally efficient and furthermore, we will see that it is asymptotically optimal in the next subsection. The effectiveness of our stopping rule is shown in Proposition 15, which ensures that provided the algorithm stops within a finite time, the probability of recommending an incorrect partition is no more than δ .

Proposition 15 *The stopping rule of BOC (Algorithm 1) ensures that*

$$\Pr(\tau_\delta < \infty, c^{\text{out}} \neq c) \leq \delta.$$

The proof of Proposition 15 is deferred to Appendix D.3. To confirm that our method BOC is indeed a δ -PAC online clustering algorithm, it remains to show it terminates within a finite time almost surely.

Remark 16 *Our stopping rule is not only applicable to our sampling rule and the estimates of (c, \mathcal{U}) produced by the K-MEANS-MAXIMIN subroutine. In fact, it also applies to any other sampling rule and the estimates by any method of estimation. In Section 6, we will experimentally compare different sampling rules with the same stopping rule.*

Remark 17 *Vaidhiyan and Sundaresan (2017) proposed a modified GLR, where the likelihood function in the numerator of the GLR statistic is replaced by an averaged likelihood function with respect to an artificial prior probability distribution, for the odd arm identification task. We remark that this method is also not practical for the online clustering task since it requires to compute one corresponding modified GLR for each possible partition. However, the total number of possible partitions is enormous (namely the Stirling number of the second kind).*

5.4 Sample Complexity Analysis

In this subsection, we analyze the correctness and sample complexity of our algorithm BOC (Algorithm 1). Proposition 18 verifies that BOC terminates within a finite time almost surely. Together with Proposition 15, it shows BOC is indeed a δ -PAC online clustering algorithm. Specifically, for any instance (c, \mathcal{U}) , BOC recommends a correct partition c^{out} based on the noisy measurements on the arm set with a probability of at least $1 - \delta$.

As shown in Proposition 18 and Theorem 19 respectively, the sample complexity of BOC asymptotically matches the instance-dependent lower bound presented in Section 4, both almost surely and in expectation, as the confidence level δ tends to zero. Therefore, BOC provably achieves asymptotic optimality in terms of the expected sample complexity and, at the same time, is also computationally efficient in terms of its sampling, stopping and recommendation rules. Thus, it achieves the best of both worlds. Refer to Appendices D.4 and D.5 for the proofs of Proposition 18 and Theorem 19, respectively.

Proposition 18 *For any instance (c, \mathcal{U}) , Algorithm 1 ensures that*

$$\Pr(\tau_\delta < \infty) = 1$$

and

$$\Pr\left(\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\log(1/\delta)} \leq D^*(c, \mathcal{U})\right) = 1.$$

Theorem 19 *For any instance (c, \mathcal{U}) , Algorithm 1 ensures that*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)} \leq D^*(c, \mathcal{U}).$$

6. Numerical Experiments

In this section, we study the empirical performance of our algorithm BOC and compare it with two baselines, namely UNIFORM and ORACLE. Since our stopping rule is the only computationally tractable one available and the recommendation rule is embedded in either the sampling rule or the stopping rule, we focus on evaluating the efficacy of our sampling rule in terms of the time it takes for the algorithm(s) to stop. As we discussed in Remark 16, any sampling rule and recommendation rule can be combined with our stopping rule. Therefore, for the sake of fairness in comparison, UNIFORM and ORACLE only differ from BOC in the sampling rules and retain the other frameworks. In particular, UNIFORM samples the M arms in a simple round-robin fashion while ORACLE samples the arms based on the optimal oracle sampling rule λ^* (i.e., the estimate $\lambda^*(t)$ in Line 8 of Algorithm 1 is replaced by λ^* , which is calculated with the unknown true pair (c, \mathcal{U})). In each setting, the reported sample complexities of different methods are averaged over 256 independent trials and the corresponding standard deviations are also shown as error bars or directly in the table. Finally, we mention that the partitions we learn in all our experiments are always correct.


6.1 Synthetic Dataset: Verifying the Asymptotic Optimality of BOC

To study the asymptotic behavior of the expected sample complexities of different methods, we construct three synthetic instances with varying difficulty levels, where $K = 4$, $M = 11$ and $d = 3$. The partitions and the first three cluster centers of all the three instances are the same, while their fourth cluster centers vary. In particular, the three instances can be expressed as follows:


$$\left\{ \begin{array}{l} c = [1, 1, 2, 2, 2, 2, 3, 3, 3, 4, 4] \\ \mu(1) = [0, 0, 0]^\top \\ \mu(2) = [0, 10, 0]^\top \\ \mu(3) = [0, 0, 10]^\top \\ \mu(4) = \begin{cases} [5, 0, 0]^\top & \text{for the } \textit{easy} \text{ instance,} \\ [1, 0, 0]^\top & \text{for the } \textit{moderate} \text{ instance,} \\ [0.5, 0, 0]^\top & \text{for the } \textit{challenging} \text{ instance.} \end{cases} \end{array} \right. \quad (9)$$

Moreover, motivated by Garivier and Kaufmann (2016), we also consider using a heuristic threshold function $\tilde{\beta}(\delta, t) := \log((1 + \log(t))^d / \delta)$, which is an approximation to the original threshold function $\beta(\delta, t)$ in Equation (8). Although no theoretical guarantee is available when we use $\tilde{\beta}(\delta, t)$, it seems practical (and even conservative) in view of the empirical error probabilities (which are always zero) for large δ (e.g., $\delta \geq 10^{-1}$).


The experimental results of the different methods with the two kinds of threshold functions for different confidence levels δ are presented in Figure 3. To better demonstrate the asymptotic behavior, we plot the empirical averaged sample complexities of the three methods as well as the instance-dependent lower bound of the expected sample complexity (see Theorem 3) with respect to $\log(1/\delta)$ in the each sub-figure. From Figure 3, we have the following observations:




(a) The easy instance with $\beta(\delta, t)$.




(b) The easy instance with $\tilde{\beta}(\delta, t)$.




(c) The moderate instance with $\beta(\delta, t)$.



(d) The moderate instance with $\tilde{\beta}(\delta, t)$.



(e) The challenging instance with $\beta(\delta, t)$.



(f) The challenging instance with $\tilde{\beta}(\delta, t)$.

Figure 3: The empirical averaged sample complexities of the different methods with the two kinds of threshold functions for different confidence levels δ on the synthetic dataset.

- Although our algorithm BOC does not require the optimal oracle sampling rule λ^* , and instead approximates it on the fly, the curves of BOC and ORACLE in the each sub-figure are almost completely overlapping, which suggests the proportion of arm pulls for BOC converges to the distribution λ^* very quickly.
- The main observation is that as δ decreases (or equivalently, as $\log(1/\delta)$ increases), the slope of the curve corresponding to the BOC algorithm in the each sub-figure is almost equal to the slope of the lower bound, which is exactly equal to $D^*(c, \mathcal{U})$. However, the slope of the curve corresponding to UNIFORM is consistently larger than that of the lower bound, exemplifying its suboptimality. This suggests that the expected sample complexity of our algorithm BOC matches the instance-dependent lower bound asymptotically, corroborating our theoretical results (see the lower bound and upper bound in Theorems 3 and 19, respectively).
- There are unavoidable gaps between the lower bound and BOC (or, almost equivalently, ORACLE), which is not an unexpected phenomenon. Although we have shown that the sampling rule resulting in λ^* is asymptotically optimal, we have to utilize the stopping rule to evaluate the quality of the final recommendation (i.e., whether the recommended partition has an error probability of at most δ). In addition, our bounds are only guaranteed to be tight *asymptotically* as $\delta \rightarrow 0$.
- Comparing the results for the same instance with the two different threshold functions, it can be seen that the heuristic one $\tilde{\beta}(\delta, t)$ leads to lower sample complexities, especially for large δ (i.e., small $\log(1/\delta)$). Even though there are no theoretical guarantees when $\tilde{\beta}(\delta, t)$ is used in our stopping rule, this threshold appears to work well empirically.

6.2 Real-world Datasets: Confirming the Non-asymptotic Superiority of BOC

To complement the experiments on synthetic data and to verify that BOC also excels in the *non-asymptotic* regime, we conduct experiments on the Iris and Yeast datasets (Dua and Graff, 2017), both of which are ubiquitous in *offline* clustering and classification tasks. Here, we perform a novel task—*online* clustering with *bandit* feedback. In the Iris dataset, the number of clusters $K = 3$, the number of arms $M = 150$, and the dimension $d = 4$, while in the Yeast dataset, $K = 10$, $M = 1484$, and $d = 8$. Note that the total number of partitions grows asymptotically as $K^M/K!$ (i.e., approximately 10^{70} or 10^{1477} in the Iris and Yeast datasets, respectively); hence, it is impractical to exhaustively enumerate over all partitions in these datasets. We emphasize that BOC succeeds in circumventing the need to solve any NP-hard optimization problem as a subroutine in the online clustering task. To adapt these datasets to be amenable to online clustering tasks, we choose each cluster center to be the mean of the original data points of the arms in it, and then rescale the centers so that the hardness parameter $D^*(c, \mathcal{U})$ is equal to 2 for both datasets. Since the performances of BOC and ORACLE are similar (as observed in Section 6.1) and the heuristic threshold function $\tilde{\beta}(\delta, t)$ generally achieves lower sample complexities (compared to when $\beta(t, \delta)$ is used), we only present the results of the two methods (namely BOC and UNIFORM) with $\tilde{\beta}(\delta, t)$ on the two real datasets.

δ	Iris Dataset		Yeast Dataset	
	BOC	UNIFORM	BOC	UNIFORM
10^{-1}	886.1 ± 55.9	1176.4 ± 69.2	14430.5 ± 371.3	19536.0 ± 530.4
10^{-2}	922.5 ± 69.4	1208.7 ± 64.1	14531.2 ± 273.4	19697.4 ± 636.2
10^{-3}	954.6 ± 80.0	1244.9 ± 72.7	14589.5 ± 175.2	19997.4 ± 718.7
10^{-4}	993.8 ± 87.3	1286.2 ± 71.9	14631.5 ± 101.6	20220.3 ± 679.2
10^{-5}	1026.9 ± 84.8	1306.0 ± 72.0	14639.6 ± 150.0	20467.7 ± 591.0
10^{-6}	1059.0 ± 68.4	1335.7 ± 67.2	14686.3 ± 218.1	20564.5 ± 499.3
10^{-7}	1075.8 ± 62.4	1368.5 ± 65.5	14723.8 ± 270.3	20686.5 ± 338.6
10^{-8}	1090.0 ± 56.9	1389.4 ± 71.7	14797.3 ± 348.5	20733.1 ± 240.6
10^{-9}	1104.8 ± 50.4	1415.6 ± 77.1	14844.6 ± 385.7	20762.0 ± 132.4
10^{-10}	1120.2 ± 48.2	1447.0 ± 73.3	14977.1 ± 445.0	20766.8 ± 107.1

Table 1: The averaged empirical sample complexities of BOC and UNIFORM with the heuristic threshold function $\tilde{\beta}(\delta, t)$ for different confidence levels δ on the real-world datasets.

The sample complexities for different confidence levels δ are presented in Table 1.³ From the table, we see that BOC significantly outperforms the non-adaptive baseline method UNIFORM for all δ in terms of sample complexities. This demonstrates that BOC is able to effectively learn the clusters in an online manner given bandit feedback.

In Appendix E, we consider a dataset with a much higher dimensionality d , namely, the MNIST dataset (LeCun et al., 1998). From Table 3 therein, we observe that the same conclusions apply.

7. Conclusion

In this paper, we proposed a novel online clustering with bandit feedback framework in which there is a set of arms that can be clustered into non-overlapping groups, and at each time, one arm is pulled, and a sample from the distribution it is associated with is observed. We proposed and analyzed BANDIT ONLINE CLUSTERING (or BOC) that, as discussed in Section 5.3, overcomes some critical computational limitations that a standard and natural GLR statistic suffers from due to the combinatorial search space of partitions. In addition to its computational efficiency, we proved that BOC is asymptotically optimal in the sense that it attains an instance-dependent information-theoretic lower bound as the confidence level δ tends to zero.

There are some limitations of the current model and theoretical contributions that serve as fertile avenues for future research. Firstly, in real-world applications such as recommendation systems and online market segmentation, it is often the case that the absolute correct

3. The instance-dependent lower bound is not presented in Table 1 since it is not informative for these two real-world datasets in the non-asymptotic regime. For instance, even when the confidence level δ is equal to 10^{-10} , the lower bound is only $D^*(c, \mathcal{U}) \cdot \log(1/\delta) \approx 46$, which is much smaller than the total number of arms (i.e., 150 or 1484 in the Iris and Yeast datasets, respectively).

clustering does not have to be found; an approximate clustering, with the advantage of further computational reductions, is usually sufficient. Developing computationally efficient and statistically optimal algorithms that allow for some distortion from the optimal clustering is thus of practical and theoretical importance. Secondly, our results are asymptotic in nature; they are only tight when the confidence level δ tends to zero. As we have seen in Section 6, this results in a gap between the upper bound and the actual performance of BOC when δ is not vanishingly small. It would thus be instructive to develop *non-asymptotic* or *refined asymptotic* bounds, perhaps by leveraging the “second-order” results in Malyutov and Tsitovich (2001) and Li and Tan (2020). Thirdly, our methodology can be generalized to the situation that the distributions of the observations are in the *multivariate exponential families*, whereas it requires efforts to preserve our computational efficiency. Finally, it is worth developing bandit feedback models and algorithms for other generalizations of clustering, such as hierarchical clustering, fuzzy or soft clustering, or community detection on graphs (Abbe, 2017).

Acknowledgments

We express our sincere gratitude to the reviewers and action editor for their meticulous reading and insightful comments. We are also thankful for the valuable discussions with Karthik Periyapattana Narayanaprasad and Yunlong Hou.

This research work is funded by the Singapore Ministry of Education Academic Research Fund (AcRF) Tier 2 under grant number A-8000423-00-00 and the Singapore Ministry of Education AcRF Tier 1 under grant number A-8000980-00-00.

Appendix A. Auxiliary Lemmas

Lemma 20 (The Maximum Theorem (Berge, 1963; Sundaram, 1996)) *Let $f : S \times \Theta \rightarrow \mathbb{R}$ be a continuous function, and $\mathcal{D} : \Theta \rightrightarrows S$ be a compact-valued continuous correspondence. Let $f^* : \Theta \rightarrow \mathbb{R}$ and $D^* : \Theta \rightrightarrows S$ be defined by*

$$f^*(\theta) = \max\{f(x, \theta) : x \in \mathcal{D}(\theta)\}$$

and

$$D^*(\theta) = \arg \max\{f(x, \theta) : x \in \mathcal{D}(\theta)\} = \{x \in \mathcal{D}(\theta) : f(x, \theta) = f^*(\theta)\}.$$

Then f^ is a continuous function on Θ , and D^* is a compact-valued, upper hemicontinuous⁴ correspondence on Θ .*

Lemma 21 (Sundaram (1996, Theorem 9.12)) *A single-valued correspondence that is hemicontinuous (whether upper or lower hemicontinuous) is continuous when viewed as a function. Conversely, every continuous function, when viewed as a single-valued correspondence, is both upper and lower hemicontinuous.*

Lemma 22 *Let λ^* denote the oracle optimal sampling rule of the instance (c, \mathcal{U}) . If there exists $\epsilon > 0$ and $t_0(\epsilon)$ such that*

$$\sup_{t \geq t_0(\epsilon)} \max_{m \in [M]} |\lambda_m^*(t) - \lambda_m^*| \leq \epsilon,$$

then there exists $t_1(\epsilon) \geq t_0(\epsilon)$ such that

$$\sup_{t \geq t_1(\epsilon)} \max_{m \in [M]} \left| \frac{N_m(t)}{t} - \lambda_m^* \right| \leq 3(M-1)\epsilon.$$

Furthermore, a valid choice of $t_1(\epsilon)$ is $\frac{t_0(\epsilon)}{\epsilon^3}$.

Proof Our sampling rule satisfies the assumptions of Garivier and Kaufmann (2016, Lemma 17), and the existence of $t_1(\epsilon)$ follows. The choice of $t_1(\epsilon)$ can be obtained by working through the proof of the same lemma. ■

4. *Hemicontinuity* of a correspondence (resp. the adjective *hemicontinuous*) is also termed as *semicontinuity* (resp. *semicontinuous*) in some books, e.g., Sundaram (1996).

Lemma 23 For any $x, y \in \mathbb{R}^+$,

$$\left((-\sqrt{x} + \sqrt{y})_+ \right)^2 = \max_{\alpha \geq 0} \left(-\alpha x + \frac{\alpha}{\alpha + 1} y \right).$$

Proof Notice that

$$-\alpha x + \frac{\alpha}{\alpha + 1} y = -(1 + \alpha)x - \frac{1}{1 + \alpha} y + (x + y).$$

If $y \geq x$, then

$$\begin{aligned} \max_{\alpha \geq 0} \left(-\alpha x + \frac{\alpha}{\alpha + 1} y \right) &= \left(-(1 + \alpha)x - \frac{1}{1 + \alpha} y \right) \Big|_{\alpha = \sqrt{\frac{y}{x}} - 1} + (x + y) \\ &= -2\sqrt{xy} + (x + y) \\ &= \left((-\sqrt{x} + \sqrt{y})_+ \right)^2. \end{aligned}$$

If $y < x$, then

$$\begin{aligned} \max_{\alpha \geq 0} \left(-\alpha x + \frac{\alpha}{\alpha + 1} y \right) &= \left(-(1 + \alpha)x - \frac{1}{1 + \alpha} y \right) \Big|_{\alpha = 0} + (x + y) \\ &= 0 \\ &= \left((-\sqrt{x} + \sqrt{y})_+ \right)^2. \end{aligned}$$

■

Lemma 24 For any $x, y \in \mathbb{R}^d$ and $\alpha \geq 0$,

$$-\alpha \|x\|^2 + \frac{\alpha}{\alpha + 1} \|y\|^2 \leq \|x - y\|^2.$$

Proof Notice that the above inequality is equivalent to

$$(1 + \alpha) \|x\|^2 - 2x \cdot y + \frac{1}{\alpha + 1} \|y\|^2 \geq 0,$$

which is equivalent to

$$\left\| \sqrt{1 + \alpha} \cdot x - \frac{1}{\sqrt{\alpha + 1}} \cdot y \right\|^2 \geq 0.$$

Thus, the result obviously holds.

■

Lemma 25 *It holds that*

$$\Pr \left(\exists t \in \mathbb{N} : \frac{1}{2} \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu(c_m)\|^2 \geq \beta(\delta, t) \right) \leq \delta$$

where $\beta(\delta, t)$ is defined in Equation (8).

Proof In this proof, we use $[x]_i$ to denote the i^{th} component of the vector $x \in \mathbb{R}^d$.

Recall that at each time t , the agent selects an arm A_t and observes $X_t = \mu(c_{A_t}) + \eta_t$, where η_t follows the standard d -dimensional Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Since each individual component of η_t independently follows the standard univariate Gaussian distribution $\mathcal{N}(0, 1)$, one arbitrary arm can be treated as d independent sub-arms. Equivalently, the agent selects a group of d sub-arms and observes

$$[X_t]_i = [\mu(c_{A_t})]_i + [\eta_t]_i$$

for all $i \in [d]$.

Note that there are Md sub-arms in total. Then Lemma 25 follows from the concentration inequality for the empirical means of sub-arms (Kaufmann and Koolen, 2021, Theorem 9). \blacksquare

Lemma 26 (Adapted from Garivier and Kaufmann (2016, Lemma 18)) *For any two constants $a > 0$ and $b \in \mathbb{R}$ such that $b + \log\left(\frac{1}{a}\right) > 0$,*

$$x = \frac{1}{a} \left(b + \log\left(\frac{e}{a}\right) + \log\left(b + \log\left(\frac{1}{a}\right)\right) \right)$$

satisfies $ax \geq \log(x) + b$.

Appendix B. Proofs of Section 4

B.1 Proof of Theorem 3

Proof For fixed $\delta \in (0, 1)$ and instance (c, \mathcal{U}) , consider any δ -PAC online clustering algorithm.

We will see in Proposition 7 that $D^*(c, \mathcal{U})$ is finite so the situation that $\mathbb{E}[\tau_\delta]$ is infinite is trivial. Henceforth, we assume that $\mathbb{E}[\tau_\delta]$ is finite.

For any arm $m \in [M]$, let $N_m(t)$ denote the number of pulls of arm m up to time t . Consider an arbitrary instance (c', \mathcal{U}') in $\text{Alt}(c)$. By applying the *transportation* inequality (Kaufmann et al., 2016, Lemma 1) and the KL-divergence for the multivariate normal distribution, we have

$$\frac{1}{2} \sum_{m=1}^M \mathbb{E}[N_m(\tau_\delta)] \|\mu(c_m) - \mu'(c'_m)\|^2 \geq d_{\text{KL}}(\delta, 1 - \delta).$$

Since the above displayed inequality holds for all instances in $\text{Alt}(c)$ and $[\mathbb{E}[N_1(\tau_\delta)], \mathbb{E}[N_2(\tau_\delta)], \dots, \mathbb{E}[N_M(\tau_\delta)]]^\top / \mathbb{E}[\tau_\delta]$ forms a probability distribution in \mathcal{P}_M , we obtain

$$\begin{aligned}
 d_{\text{KL}}(\delta, 1 - \delta) &\leq \frac{1}{2} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \mathbb{E}[N_m(\tau_\delta)] \|\mu(c_m) - \mu'(c'_m)\|^2 \\
 &= \frac{1}{2} \mathbb{E}[\tau_\delta] \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \frac{\mathbb{E}[N_m(\tau_\delta)]}{\mathbb{E}[\tau_\delta]} \|\mu(c_m) - \mu'(c'_m)\|^2 \\
 &\leq \frac{1}{2} \mathbb{E}[\tau_\delta] \sup_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \\
 &= \mathbb{E}[\tau_\delta] D^*(c, \mathcal{U})^{-1}.
 \end{aligned}$$

Since $\lim_{\delta \rightarrow 0} d_{\text{KL}}(\delta, 1 - \delta) / \log(1/\delta) = 1$, letting $\delta \rightarrow 0$ yields

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)} \geq D^*(c, \mathcal{U})$$

as desired. ■

B.2 Proof of Lemma 4

Proof For any fixed $\lambda \in \mathcal{P}_M$ and (c', \mathcal{U}') , let

$$\text{Dist}(c', \mathcal{U}') := \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2.$$

To prove Lemma 4, we only need to show for any instance $(c^\dagger, \mathcal{U}^\dagger) \in \text{Alt}(c)$ such that $d_{\text{H}}(c^\dagger, c) > 1$, there exists another instance $(c^*, \mathcal{U}^*) \in \text{Alt}(c)$ such that $d_{\text{H}}(c^*, c) = 1$ and $\text{Dist}(c^*, \mathcal{U}^*) \leq \text{Dist}(c^\dagger, \mathcal{U}^\dagger)$. The proof consists of four steps and Figure 4 serves as an illustration to help understand the various constructions.

Step 1 (Permute the partition of the given instance). To construct a new instance, we construct a permutation σ on $[K]$ such that

$$\left\{ \begin{array}{l}
 \sigma^{-1}(1) = \arg \min_{k \in [K]} \|\mu^\dagger(k) - \mu(1)\|^2 \\
 \sigma^{-1}(2) = \arg \min_{k \in [K] \setminus \{\sigma^{-1}(1)\}} \|\mu^\dagger(k) - \mu(2)\|^2 \\
 \sigma^{-1}(3) = \arg \min_{k \in [K] \setminus \{\sigma^{-1}(1), \sigma^{-1}(2)\}} \|\mu^\dagger(k) - \mu(3)\|^2 \\
 \dots \\
 \sigma^{-1}(K-1) = \arg \min_{k \in [K] \setminus \{\sigma^{-1}(1), \sigma^{-1}(2), \dots, \sigma^{-1}(K-2)\}} \|\mu^\dagger(k) - \mu(K-1)\|^2.
 \end{array} \right. \quad (10)$$




Figure 4: An illustration of the proof of Lemma 4. Figure 4 illustrates the construction of the sequence of instances when the number of clusters $K = 3$ and the number of arms (or items) $M = 9$. Each pair of items connected by a double arrow represents the cluster indices of one arm in the true partition c and one of the newly constructed partitions $c^{(\bar{k})}$ for $\bar{k} = 0, 1, 2, 3$. After the application of the permutation σ as defined in (10) in Step 1, each of the arms can have any cluster index. However, due to the desirable property of the permutation as stated in (11), in Step 2, we are able to construct a new partition $c^{(1)}$ such that for any arm, its cluster index in $c^{(1)}$ is not larger than that of c . Next, we modify the new partition from right to left in Step 3 (see Equations (13) and (14)) and we finally return to a partition that is identical to c .

Let $(c^{(0)}, \mathcal{U}^{(0)}) = \sigma(c^\dagger, \mathcal{U}^\dagger)$ and hence we have

$$\begin{cases} 1 = \arg \min_{1 \leq k \leq K} \|\mu^{(0)}(k) - \mu(1)\|^2 \\ 2 = \arg \min_{2 \leq k \leq K} \|\mu^{(0)}(k) - \mu(2)\|^2 \\ 3 = \arg \min_{3 \leq k \leq K} \|\mu^{(0)}(k) - \mu(3)\|^2 \\ \dots \\ K-1 = \arg \min_{K-1 \leq k \leq K} \|\mu^{(0)}(k) - \mu(K-1)\|^2. \end{cases} \quad (11)$$

Obviously, $d_{\text{H}}(c, c^{(0)}) \geq 1$, otherwise $(c^\dagger, \mathcal{U}^\dagger) \notin \text{Alt}(c)$.

Step 2 (Update the clustering). Now we construct another instance $(c^{(1)}, \mathcal{U}^{(0)})$, in which $c_m^{(1)} = \min(c_m, c_m^{(0)})$ for all $m \in [M]$. It holds that

$$\begin{aligned}
 & \text{Dist}(c^{(0)}, \mathcal{U}^{(0)}) \\
 &= \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m^{(0)})\|^2 \\
 &= \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m^{(0)})\|^2 \mathbb{1}\{c_m \leq c_m^{(0)}\} + \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m^{(0)})\|^2 \mathbb{1}\{c_m > c_m^{(0)}\} \\
 &\geq \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m)\|^2 \mathbb{1}\{c_m \leq c_m^{(0)}\} + \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m^{(0)})\|^2 \mathbb{1}\{c_m > c_m^{(0)}\} \\
 &= \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu^{(0)}(c_m^{(1)})\|^2 \\
 &= \text{Dist}(c^{(1)}, \mathcal{U}^{(0)})
 \end{aligned}$$

where the inequality is due to our construction of $(c^{(0)}, \mathcal{U}^{(0)})$ in Equation (11).

Step 3 (Construct a sequence of instances). For any (c', \mathcal{U}') and $k, k' \in [K]$, let

$$w_{k,k'}(c') := \sum_{m=1}^M \lambda_m \mathbb{1}\{c_m = k, c'_m = k'\}.$$

Then

$$\text{Dist}(c', \mathcal{U}') = \sum_{k=1}^K \sum_{k'=1}^K w_{k,k'}(c') \|\mu(k) - \mu'(k')\|^2.$$

Moreover, the minimization problem $\min_{\mathcal{U}'} \text{Dist}(c', \mathcal{U}')$ has a unique solution \mathcal{U}'^* which is defined by its components as

$$\mu'^*(k') = \frac{\sum_{k=1}^K w_{k,k'}(c') \mu(k)}{\sum_{k=1}^K w_{k,k'}(c')} \quad (12)$$

for all $k' \in [K]$.

Consider an instance $(c^{(1)}, \mathcal{U}^{(1)})$, in which $\mathcal{U}^{(1)} = \arg \min_{\mathcal{U}'} \text{Dist}(c^{(1)}, \mathcal{U}')$. Clearly, it holds that $\text{Dist}(c^{(1)}, \mathcal{U}^{(0)}) \geq \text{Dist}(c^{(1)}, \mathcal{U}^{(1)})$. From the construction of $c^{(1)}$, we also know that $w_{k,K}(c^{(1)}) = 0$ for any $k \neq K$. Therefore, by (12), $\mu^{(1)}(K) = \mu(K)$.

Then consider another instance $(c^{(2)}, \mathcal{U}^{(1)})$, where

$$c_m^{(2)} = \begin{cases} K & \text{if } c_m = K \\ c_m^{(1)} & \text{otherwise.} \end{cases} \quad (13)$$

Since $\mu^{(1)}(K) = \mu(K)$, $\text{Dist}(c^{(1)}, \mathcal{U}^{(1)}) \geq \text{Dist}(c^{(2)}, \mathcal{U}^{(1)})$. Next, we construct $(c^{(2)}, \mathcal{U}^{(2)})$, in which $\mathcal{U}^{(2)} = \arg \min_{\mathcal{U}'} \text{Dist}(c^{(2)}, \mathcal{U}')$. From the construction of $c^{(2)}$, we also know that

$w_{k,K-1}(c^{(2)}) = 0$ for any $k \neq K-1$. Therefore, $\mu^{(2)}(K-1) = \mu(K-1)$ and $\text{Dist}(c^{(2)}, \mathcal{U}^{(2)}) \geq \text{Dist}(c^{(3)}, \mathcal{U}^{(2)})$ where

$$c_m^{(3)} = \begin{cases} K-1 & \text{if } c_m = K-1 \\ c_m^{(2)} & \text{otherwise.} \end{cases} \quad (14)$$

Following the same method, we can then construct a sequence of instances $\{(c^{(3)}, \mathcal{U}^{(3)}), (c^{(4)}, \mathcal{U}^{(3)}), \dots, (c^{(K)}, \mathcal{U}^{(K)})\}$. Note that eventually $(c^{(K)}, \mathcal{U}^{(K)}) = (c, \mathcal{U})$ since the cluster indices of $c^{(K)}$ and c are the same and $\mathcal{U}^{(K)} = \arg \min_{\mathcal{U}'} \text{Dist}(c^{(K)}, \mathcal{U}')$.

Step 4 (Find the desired instance within the constructed sequence). Consider the whole sequence $\{(c^{(0)}, \mathcal{U}^{(0)}), (c^{(1)}, \mathcal{U}^{(0)}), (c^{(1)}, \mathcal{U}^{(1)}), (c^{(2)}, \mathcal{U}^{(1)}), \dots, (c^{(K)}, \mathcal{U}^{(K)})\}$. Both the distance function $\text{Dist}(\cdot, \cdot)$ and the Hamming distance $d_{\text{H}}(c, \cdot)$ are non-increasing. Furthermore, $d_{\text{H}}(c, c^{(K)}) = 0$.

If $d_{\text{H}}(c, c^{(0)}) = 1$, we can simply choose $(c^*, \mathcal{U}^*) = (c^{(0)}, \mathcal{U}^{(0)})$.

If $d_{\text{H}}(c, c^{(0)}) > 1$, from the sequence of instances, we can select the first instance whose partition is exactly the same as c , denoted as $(c^{(\bar{k})}, \mathcal{U}^{(\bar{k}-1)})$ ($1 \leq \bar{k} \leq K$), i.e., $c^{(\bar{k})} = c$. We also know $(c^{(\bar{k}-1)}, \mathcal{U}^{(\bar{k}-1)})$ satisfies that $d_{\text{H}}(c, c^{(\bar{k}-1)}) \geq 1$.

Our construction of the sequence of instances is “consecutive” in the sense that we can modify the cluster indices of $c^{(\bar{k}-1)}$ one by one until we get exactly the partition $c^{(\bar{k})}$. Therefore, we can construct another sequence of instances $\{(c^{(\bar{k}-1)}, \mathcal{U}^{(\bar{k}-1)}), (c^{(\bar{k}'-1)}, \mathcal{U}^{(\bar{k}-1)}), (c^{(\bar{k}''-1)}, \mathcal{U}^{(\bar{k}-1)}), \dots, (c^{(\bar{k})}, \mathcal{U}^{(\bar{k}-1)})\}$, where the Hamming distance $d_{\text{H}}(c, \cdot)$ is strictly decreasing by 1 in each step in this sequence. Besides, the distance function $\text{Dist}(\cdot, \cdot)$ is also non-increasing. As a result, we can always find an instance (c^*, \mathcal{U}^*) in this sequence such that $d_{\text{H}}(c^*, c) = 1$. \blacksquare

B.3 Proof of Proposition 5

Proof Due to Lemma 4, we only need to consider the set of alternative instances whose partitions have a Hamming distance of 1 from the given partition c , i.e., $\{(c', \mathcal{U}') \in \text{Alt}(c) : d_{\text{H}}(c', c) = 1\}$.

If $\lambda \in \mathcal{P}_M^+$, suppose that c and c' only differ in the label of arm m , which is changed from c_m to k' . Since c' is a valid partition, we have $n(c_m) > 1$. To minimize the objective function, for any $k \in [K] \setminus \{k'\}$, $\mu'(k)$ can be exactly set to be $\mu(k)$. Therefore, the objective function can be simplified to

$$\lambda_m \|\mu(c_m) - \mu'(k')\|^2 + w(k') \|\mu(k') - \mu'(k')\|^2,$$

which shows the mean vector of cluster k' should be chosen as $\mu'(k') = \frac{\lambda_m \mu(c_m) + w(k') \mu(k')}{\lambda_m + w(k')}$, a weighted sum of $\mu(c_m)$ and $\mu(k')$.

Altogether, the optimal value can be derived as follows:

$$\begin{aligned}
 & \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \\
 = & \inf_{\substack{(c', \mathcal{U}') \in \text{Alt}(c): \\ d_{\text{H}}(c', c) = 1}} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \\
 = & \inf_{\substack{m \in [M], k' \in [K]: \\ n(c_m) > 1, k' \neq c_m}} \frac{\lambda_m w(k')}{\lambda_m + w(k')} \|\mu(c_m) - \mu(k')\|^2 \tag{15}
 \end{aligned}$$

$$\begin{aligned}
 = & \min_{\substack{m \in [M], k' \in [K]: \\ n(c_m) > 1, k' \neq c_m}} \frac{\lambda_m w(k')}{\lambda_m + w(k')} \|\mu(c_m) - \mu(k')\|^2 \tag{16}
 \end{aligned}$$

$$\begin{aligned}
 = & \min_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \frac{\bar{w}(k) w(k')}{\bar{w}(k) + w(k')} \|\mu(k) - \mu(k')\|^2 \tag{17}
 \end{aligned}$$

where Equation (15) follows from our choice of the mean vectors above and Equation (17) follows from the fact that for any $y, z > 0$, $f(x) = \frac{xy}{x+z}$ is increasing on $(0, \infty)$. As a consequence, the infimum in Problem (Δ) can be replaced with a minimum since the infimum can be attained by some $(c', \mathcal{U}') \in \text{Alt}(c)$.

If $\lambda \in \mathcal{P}_M \setminus \mathcal{P}_M^+$, suppose that $\lambda_m = 0$ for some $m \in [M]$. Notice that the objective function is always non-negative.

We consider two situations. First, if $n(c_m) > 1$, we will construct an instance $(c', \mathcal{U}') \in \text{Alt}(c)$ such that the objective function is zero. In particular, we can change the label of the arm m from c_m to any $k \neq c_m$, and keep all the mean vectors the same. Then the objective function for the new instance is exactly zero, which shows

$$\inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 = 0.$$

Second, if $n(c_m) = 1$, we will construct an instance $(c', \mathcal{U}') \in \text{Alt}(c)$ such that the objective function is arbitrarily close to zero. In particular, we can change the label of any other arm m' (subject to $n(c_{m'}) > 1$) from $c_{m'}$ to c_m . Thus the objective function can be simplified to

$$\lambda_{m'} \|\mu(c_{m'}) - \mu'(c_m)\|^2.$$

For the mean vector of c_m , note that we cannot set $\mu'(c_m)$ to be exactly $\mu(c_{m'})$ otherwise $(c', \mathcal{U}') \notin \text{Alt}(c)$. However, we can let $\mu'(c_m)$ be arbitrarily close to $\mu(c_{m'})$, which yields

$$\inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 = 0.$$

This completes the proof of Proposition 5. ■

Remark 27 In fact, from the above proof, we know more about the optimal solution if $\lambda \in \mathcal{P}_M^+$. Let

$$(k^*, k'^*) := \arg \min_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \frac{\bar{w}(k)w(k')}{\bar{w}(k) + w(k')} \|\mu(k) - \mu(k')\|^2.$$

Then there exists (c^*, \mathcal{U}^*) , where

$$c_m^* = \begin{cases} k'^* & \text{if } m = \arg \min_{m \in [M]: c_m = k^*} \lambda_m \\ c_m & \text{otherwise} \end{cases}$$

and

$$\mu^*(k) = \begin{cases} \frac{\bar{w}(k^*)\mu(k^*) + w(k'^*)\mu(k'^*)}{\bar{w}(k^*) + w(k'^*)} & \text{if } m = k'^* \\ \mu(k) & \text{otherwise} \end{cases}$$

such that

$$(c^*, \mathcal{U}^*) \in \arg \min_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2.$$

B.4 Proof of Proposition 6

Proof First, we define $f_1 : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ as

$$f_1(x, y) := \begin{cases} \frac{xy}{x+y} & \text{if } x + y \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

which is continuous on $\mathbb{R}^+ \times \mathbb{R}^+$.

Next, we will prove that for all $(\lambda, \mathcal{U}) \in \mathcal{P}_M \times \mathcal{U}$, $g(\lambda, \mathcal{U})$ as defined in Proposition 6 can be written as

$$g(\lambda, \mathcal{U}) = \min_{\substack{m \in [M], k \in [K]: \\ n(c_m) > 1, k \neq c_m}} f_1 \left(\lambda_m, \sum_{\bar{m}=1}^M \lambda_{\bar{m}} \mathbb{1}\{c_{\bar{m}} = k\} \right) \|\mu(c_m) - \mu(k)\|^2. \quad (18)$$

For any $\lambda \in \mathcal{P}_M^+$, Equation (18) follows directly from Equation (16) in the proof of Proposition 5.

For any $\lambda \in \mathcal{P}_M \setminus \mathcal{P}_M^+$, it suffices to find $m \in [M], k \in [K]$ subject to $n(c_m) > 1$ and $k \neq c_m$ such that $f_1 \left(\lambda_m, \sum_{\bar{m}=1}^M \lambda_{\bar{m}} \mathbb{1}\{c_{\bar{m}} = k\} \right) = 0$. Suppose that $\lambda_{\hat{m}} = 0$ for some $\hat{m} \in [M]$ and we will consider two situations. If $n(c_{\hat{m}}) > 1$, we can simply choose $m = \hat{m}$ and any $k \neq c_{\hat{m}}$. If $n(c_{\hat{m}}) = 1$, we can choose any m subject to $n(c_m) > 1$, and $k = c_{\hat{m}}$. Altogether, Equation (18) holds for all $(\lambda, \mathcal{U}) \in \mathcal{P}_M \times \mathcal{U}$.

Finally, for any $m \in [M], k \in [K]$ such that $n(c_m) > 1$ and $k \neq c_m$, we consider the function

$$f_2(m, k; \lambda, \mathcal{U}) := f_1 \left(\lambda_m, \sum_{\bar{m}=1}^M \lambda_{\bar{m}} \mathbb{1}\{c_{\bar{m}} = k\} \right) \|\mu(c_m) - \mu(k)\|^2$$

so that

$$g(\lambda, \mathcal{U}) = \min_{\substack{m \in [M], k \in [K]: \\ n(c_m) > 1, k \neq c_m}} f_2(m, k; \lambda, \mathcal{U}).$$

Since $(\lambda, \mathcal{U}) \mapsto f_2(m, k; \lambda, \mathcal{U})$ is continuous on $\mathcal{P}_M \times \mathcal{U}$ for fixed m and k and the finite minimum operation preserves continuity, $(\lambda, \mathcal{U}) \mapsto g(\lambda, \mathcal{U})$ is also continuous on $\mathcal{P}_M \times \mathcal{U}$. \blacksquare

B.5 Proof of Proposition 7

Proof For any (c, \mathcal{U}) , $D^*(c, \mathcal{U})$ can be written as:

$$\begin{aligned} D^*(c, \mathcal{U}) &= \left\{ \frac{1}{2} \sup_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \right\}^{-1} \\ &= \left\{ \frac{1}{2} \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \right\}^{-1} \end{aligned} \quad (19)$$

$$\begin{aligned} &= \left\{ \frac{1}{2} \max_{\lambda \in \mathcal{P}_M^+} \min_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \frac{\bar{w}(k)w(k')}{\bar{w}(k) + w(k')} \|\mu(k) - \mu(k')\|^2 \right\}^{-1} \\ &= 2 \min_{\lambda \in \mathcal{P}_M^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{1}{\bar{w}(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2}. \end{aligned} \quad (20)$$

where Equation (19) follows from the continuity of $g(\lambda, \mathcal{U})$ as shown in Proposition 6 and the compactness of \mathcal{P}_M , and Equation (20) follows from Proposition 5 and the non-negativity of $D^*(c, \mathcal{U})$.

Notice that both $\{w(k)\}_{k=1}^K$ and $\{\bar{w}(k)\}_{k=1}^K$ depend on λ . For any k and fixed $w(k)$, $\bar{w}(k) = \min_{m \in [M]: c_m = k} \lambda_m$ is maximized *if and only if* for all $m \in [M]$ such that $c_m = k$, λ_m are equal. Therefore, we can solve the outer minimization on a “smaller” probability simplex \mathcal{P}_K^+ . This yields

$$D^*(c, \mathcal{U}) = 2 \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2}$$

as desired. \blacksquare

B.6 Proof of Proposition 8

Proof According to the proof of Proposition 7, there exists a bijective map (which is specified in the proof of Proposition 7 as well as the statement of Proposition 8) between the solution(s) to

$$\arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2$$

and the solution(s) to

$$\arg \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2}. \quad (21)$$

Therefore, it suffices to show the solution to (21) is unique. We will show this by contradiction as follows.

Suppose that w^* and w^{**} are different solutions to (21) and $w^*(\hat{k}) \neq w^{**}(\hat{k})$ for some $\hat{k} \in [K]$.

For any $k, k' \in [K]$ such that $n(k) > 1$ and $k' \neq k$, consider the function

$$f_1(k, k'; w, \mathcal{U}) := 2 \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2}.$$

Since $\|\mu(k) - \mu(k')\|^{-2} > 0$ by the definition of \mathcal{U} , $f_1(k, k'; w, \mathcal{U})$ is convex in w . In particular, $f_1(k, k'; w, \mathcal{U})$ is strictly convex in the pair $(w(k), w(k'))$. Besides, since the pointwise maximum operation preserves convexity,

$$f_2(w, \mathcal{U}) := \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} f_1(k, k'; w, \mathcal{U})$$

is also convex in w . Therefore, $w^{***} := \frac{1}{2}(w^* + w^{**})$ is also a solution to (21) and

$$D^*(c, \mathcal{U}) = f_2(w^*, \mathcal{U}) = f_2(w^{**}, \mathcal{U}) = f_2(w^{***}, \mathcal{U}).$$

We claim that for any $k, k' \in [K]$ such that $n(k) > 1$ and $k' \neq k$, if $\hat{k} \in \{k, k'\}$, then

$$f_1(k, k'; w^{***}, \mathcal{U}) < D^*(c, \mathcal{U}).$$

Otherwise, using the fact $w^*(\hat{k}) \neq w^{**}(\hat{k})$ and the strict convexity of $f_1(k, k'; w^{***}, \mathcal{U})$ in the pair $(w^{***}(k), w^{***}(k'))$,

$$\begin{aligned} D^*(c, \mathcal{U}) &= f_1(k, k'; w^{***}, \mathcal{U}) \\ &< \frac{1}{2} (f_1(k, k'; w^*, \mathcal{U}) + f_1(k, k'; w^{**}, \mathcal{U})) \\ &\leq \frac{1}{2} (f_2(w^*, \mathcal{U}) + f_2(w^{**}, \mathcal{U})) \\ &= D^*(c, \mathcal{U}) \end{aligned}$$

which leads to a contradiction. Therefore, our claim holds.

In view of this claim, we can identify a $w^\dagger \in \mathcal{P}_K^+$ such that $f_2(w^\dagger, \mathcal{U}) < D^*(c, \mathcal{U})$, which results in a contradiction to the optimality of $D^*(c, \mathcal{U})$. In particular, such a w^\dagger can be defined through its components as

$$w^\dagger(k) = \begin{cases} w^{***}(k) - (K-1)\epsilon & \text{if } k = \hat{k} \\ w^{***}(k) + \epsilon & \text{otherwise} \end{cases}$$

with a sufficiently small $\epsilon > 0$ such that

$$f_1(k, k'; w^\dagger, \mathcal{U}) < D^*(c, \mathcal{U})$$

for all $k, k' \in [K]$ such that $n(k) > 1$, $k' \neq k$ and $\hat{k} \in \{k, k'\}$. The existence of such a sufficiently small $\epsilon > 0$ is guaranteed by the continuity of $f_1(k, k'; w^{***}, \mathcal{U})$ in the pair $(w^{***}(k), w^{***}(k'))$.

In addition, for any $k, k' \in [K]$ such that $n(k) > 1$, $k' \neq k$ and $\hat{k} \notin \{k, k'\}$,

$$\begin{aligned} f_1(k, k'; w^\dagger, \mathcal{U}) &= 2 \left(\frac{n(k)}{w^\dagger(k)} + \frac{1}{w^\dagger(k')} \right) \|\mu(k) - \mu(k')\|^{-2} \\ &< 2 \left(\frac{n(k)}{w^{***}(k)} + \frac{1}{w^{***}(k')} \right) \|\mu(k) - \mu(k')\|^{-2} \\ &= f_1(k, k'; w^{***}, \mathcal{U}) \\ &\leq D^*(c, \mathcal{U}). \end{aligned}$$

In view of the fact that under both cases ($\hat{k} \in \{k, k'\}$ and $\hat{k} \notin \{k, k'\}$), we have $f_1(k, k'; w^\dagger, \mathcal{U}) < D^*(c, \mathcal{U})$, we conclude that

$$f_2(w^\dagger, \mathcal{U}) = \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} f_1(k, k'; w^\dagger, \mathcal{U}) < D^*(c, \mathcal{U})$$

which contradicts the fact that $D^*(c, \mathcal{U})$ is the optimal value.

Altogether, the solution to (21) is unique and hence the solution to

$$\arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2$$

is also unique. ■

B.7 Proof of Proposition 9

Proof Recalling Proposition 6,

$$g(\lambda, \mathcal{U}) = \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2$$

is continuous on $\mathcal{P}_M \times \mathcal{U}$.

Since \mathcal{P}_M does not depend on \mathcal{U} and is compact, by Lemma 20,

$$\Lambda(\mathcal{U}) = \arg \max_{\lambda \in \mathcal{P}_M} g(\lambda, \mathcal{U})$$

is an upper hemicontinuous correspondence on \mathcal{U} .

According to Proposition 8, $\Lambda(\mathcal{U})$ is single-valued. Consequently, by Lemma 21, Λ is continuous on \mathcal{U} . ■

Appendix C. On the Forced Exploration Stage of Algorithm 1

In this appendix, we study the situation that the ratio between the minimal and the maximal pairwise distances among the cluster centers is bounded, and propose a novel exploration method, which is more adaptive than the original forced exploration used in Algorithm 1. In particular, we assume that (c, \mathcal{U}) , the instance of cluster bandits to be partitioned, satisfies Assumption 1.

Assumption 1 *The ratio between the minimal and the maximal pairwise distances among the K centers of the clusters $\{\mu(k)\}_{k \in [K]}$ is lower-bounded by a known constant $r \in (0, 1]$, i.e.,*

$$\frac{\min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|}{\max_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|} \geq r.$$

We start from a quantitative result on the optimal oracle sampling rule λ^* , which is proved in Appendix C.1.

Proposition 28 *Under Assumption 1, λ^* , which is the unique solution to*

$$\arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2,$$

satisfies $\lambda_m^ \geq \frac{r^2}{2M}$ for all $m \in [M]$.*

Basically, Proposition 28 shows that the proportions of arm pulls for the optimal sampling rule λ^* are uniformly lower-bounded by a positive constant. Therefore, at each time step t , we no longer need to check whether there exists an arm whose number of pulls is not larger than a preset threshold. Once we obtain the plug-in approximation $\lambda^*(t)$ with the estimate produced by K-MEANS-MAXIMIN, we project it onto the constrained probability simplex $\mathcal{P}_M(\frac{r^2}{2M}) := \{\lambda \in [\frac{r^2}{2M}, 1]^M : \|\lambda\|_1 = 1\}$. Then instead of the original plug-in approximation $\lambda^*(t)$ which may result in under-sampling, we can track the optimal sampling rule based on this projection (which is denoted as $\tilde{\lambda}^*(t)$). In summary, the sampling rule (Line 8 of Algorithm 1) becomes

$$A_{t+1} = \arg \max_{m \in [M]} t \tilde{\lambda}_m^*(t) - N_m(t) \quad (22)$$

where

$$\tilde{\lambda}^*(t) = \text{Proj}_{\mathcal{P}_M(\frac{r^2}{2M})}(\lambda^*(t)) = \arg \min_{\lambda \in \mathcal{P}_M(\frac{r^2}{2M})} \|\lambda - \lambda^*(t)\|.$$

We refer to the improved sampling rule in Equation (22) as the linear-exploration sampling rule henceforth. Compared to the original one proposed in Section 5.2, the linear-exploration sampling rule circumvents the forced exploration stage and always maintains a linear minimal exploration rate, which ensures that the plug-in approximation $\lambda^*(t)$ converges to the oracle optimal sampling rule λ^* (almost surely). Moreover, since Proposition 28 indicates that the optimal sampling rule always lies in the constrained probability

simplex $\mathcal{P}_M(\frac{r^2}{2M})$ under Assumption 1, the projection $\tilde{\lambda}^*(t)$ also converges to λ^* (almost surely). Therefore, asymptotic optimality is preserved. Besides, one can easily see that our subsequent results on sample complexity (Proposition 18 and Theorem 19) also hold. However, since the forced exploration stage in the original sampling rule of BOC might take effect even if the plug-in approximation $\lambda^*(t)$ is close to the optimal sampling rule λ^* , the linear-exploration sampling rule may result in better empirical performance in the non-asymptotic regime, as demonstrated in the numerical results in Appendix C.2.

Finally, we remark that the linear-exploration sampling rule requires extra knowledge of the pairwise distances among the cluster centers. This is unlikely to be available in most practical applications and hence we adopt an arguably less elegant forced exploration strategy in the main text.

C.1 Proof of Proposition 28

Proof By Proposition 8, it suffices to show that w^* , which is the unique solution to

$$\arg \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2},$$

satisfies $w^*(k) \geq \frac{r^2 n(k)}{2M}$ for all $k \in [K]$.

In the following, we will show this result by contradiction.

Suppose that there exists $\hat{k} \in [K]$ such that $w^*(\hat{k}) < \frac{r^2 n(\hat{k})}{2M}$. For ease of notation, we denote

$$\begin{aligned} (\spadesuit) &:= \min_{w \in \mathcal{P}_K^+} \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w(k)} + \frac{1}{w(k')} \right) \|\mu(k) - \mu(k')\|^{-2} \\ &= \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{w^*(k)} + \frac{1}{w^*(k')} \right) \|\mu(k) - \mu(k')\|^{-2}. \end{aligned}$$

If $n(\hat{k}) = 1$, then

$$\begin{aligned} (\spadesuit) &\geq \max_{\substack{k \in [K]: \\ n(k) > 1, k \neq \hat{k}}} \left(\frac{n(k)}{w^*(k)} + \frac{1}{w^*(\hat{k})} \right) \|\mu(k) - \mu(\hat{k})\|^{-2} \\ &> \frac{1}{w^*(\hat{k})} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2} \\ &> \frac{2M}{r^2} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2}. \end{aligned}$$

If $n(\hat{k}) > 1$, then

$$\begin{aligned}
(\spadesuit) &\geq \max_{k' \in [K]: k' \neq \hat{k}} \left(\frac{n(\hat{k})}{w^*(\hat{k})} + \frac{1}{w^*(k')} \right) \|\mu(\hat{k}) - \mu(k')\|^{-2} \\
&> \frac{n(\hat{k})}{w^*(\hat{k})} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2} \\
&> \frac{2M}{r^2} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2}.
\end{aligned}$$

Therefore, in both cases, we have

$$(\spadesuit) > \frac{2M}{r^2} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2}. \quad (23)$$

On the other hand, we construct $\tilde{w} \in \mathbb{R}^K$ such that $\tilde{w}(k) = \frac{n(k)}{M}$ for all $k \in [K]$. Since $\sum_{k \in [K]} n(k) = M$, \tilde{w} is a valid element in \mathcal{P}_K^+ . Thus, under Assumption 1, it holds that

$$\begin{aligned}
(\spadesuit) &\leq \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(\frac{n(k)}{\tilde{w}(k)} + \frac{1}{\tilde{w}(k')} \right) \|\mu(k) - \mu(k')\|^{-2} \\
&= \max_{\substack{k, k' \in [K]: \\ n(k) > 1, k' \neq k}} \left(M + \frac{M}{n(k')} \right) \|\mu(k) - \mu(k')\|^{-2} \\
&\leq 2M \max_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2} \\
&\leq \frac{2M}{r^2} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|^{-2}
\end{aligned}$$

which contradicts Inequality (23).

The proof of Proposition 28 is thus completed. \blacksquare

C.2 Numerical Experiments on the Linear-Exploration Sampling Rule

In this section, we numerically compare the more aggressive linear-exploration sampling rule described in Equation (22) to the the original sampling rule proposed in Section 5.2. We experiment this on the challenging synthetic instance in Equation (9) and the Iris dataset as these datasets behave differently. For the challenging synthetic instance, we rescale the cluster centers such that its hardness parameter $D^*(c, \mathcal{U})$ is equal to that of the Iris dataset. For the linear-exploration sampling rule, the exact information of the ratio between the minimal and the maximal pairwise distances is given to the algorithm as an input.

The empirical stopping times for various (non-vanishing) failure probabilities δ are presented in Table 2. The behaviors on both datasets are different. The linear-exploration rule outperforms the original one for the rescaled challenging synthetic instance, while there are no obvious gaps between the two sampling rules for the Iris dataset. This is because for the rescaled challenging instance, there exists some arms that require a few arm pulls

δ	Rescaled Challenging Instance		Iris Dataset	
	ORIGINAL	LE	ORIGINAL	LE
10^{-1}	102.9 ± 29.5	78.9 ± 20.3	886.1 ± 55.9	886.8 ± 47.5
10^{-2}	119.0 ± 28.9	86.7 ± 21.3	922.5 ± 69.4	920.7 ± 71.8
10^{-3}	132.3 ± 29.4	95.9 ± 22.3	954.6 ± 80.0	953.4 ± 77.8
10^{-4}	142.1 ± 31.5	102.8 ± 23.6	993.8 ± 87.3	990.1 ± 88.8
10^{-5}	158.3 ± 32.7	114.7 ± 25.4	1026.9 ± 84.8	1023.6 ± 83.7
10^{-6}	165.8 ± 32.3	120.1 ± 26.5	1059.0 ± 68.4	1055.3 ± 71.8
10^{-7}	176.1 ± 31.2	128.9 ± 27.1	1075.8 ± 62.4	1079.1 ± 58.8
10^{-8}	189.1 ± 35.3	135.4 ± 26.7	1090.0 ± 56.9	1092.2 ± 54.1
10^{-9}	199.6 ± 32.8	143.3 ± 27.3	1104.8 ± 50.4	1109.6 ± 48.7
10^{-10}	209.7 ± 33.8	152.2 ± 29.6	1120.2 ± 48.2	1120.0 ± 49.7

Table 2: The averaged empirical sample complexities of BOC with the original forced-exploration sampling rule described in Algorithm 1 (and indicated by ORIGINAL in the table) or the linear-exploration Sampling Rule (indicated by LE) with the heuristic threshold function $\tilde{\beta}(\delta, t)$ for different confidence levels δ .

to learn their means sufficiently accurately. However, due to the forced exploration rule, they are pulled more often than necessary and this results in the inefficiency. On the other hand, this phenomenon is not observed for the Iris dataset and both Algorithm 1 and the linear-exploration rule work well in ensuring that the approximate proportions of arm pulls ($\lambda^*(t)$ or $\tilde{\lambda}^*(t)$) converge to the optimal one dictated by the lower bound.

Appendix D. Proofs of Section 5

D.1 Proof of Proposition 10

Proof The proof consists of three steps, using the triangle inequality frequently.

First, we show that the K arms chosen in Maximin Initialization are selected from K disjoint clusters.

Let m_1, m_2, \dots, m_K denote the arms chosen in Maximin Initialization in order. Assuming that we have already taken the empirical estimates of \bar{k} arms as the cluster centers, consider $\mathcal{A}_{\bar{k}}$, the set of the arms that shares the same true cluster index with at least one of the existing centers. For any arm $m' \in \mathcal{A}_{\bar{k}}$, the Euclidean distance to the nearest existing center can be upper bounded as follows:

$$\begin{aligned}
 \min_{1 \leq k \leq \bar{k}} \|\hat{\mu}_{m'} - \hat{\mu}(k)\| &= \min_{1 \leq k \leq \bar{k}} \|\hat{\mu}_{m'} - \hat{\mu}_{m_k}\| \\
 &\leq \min_{1 \leq k \leq \bar{k}} (\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| + \|\mu(c_{m_k}) - \hat{\mu}_{m_k}\|) \\
 &\leq \min_{1 \leq k \leq \bar{k}} \left(2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| \right) \\
 &= 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{1 \leq k \leq \bar{k}} \|\mu(c_{m'}) - \mu(c_{m_k})\|
 \end{aligned}$$

$$= 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|$$

where the last equality results from the fact that at least one of the existing centers shares the same true cluster index as that of arm m' .

On the other hand, for any arm $m' \in [M] \setminus \mathcal{A}_{\bar{k}}$, the Euclidean distance to the nearest existing center can be lower bounded as follows:

$$\begin{aligned} \min_{1 \leq k \leq \bar{k}} \|\hat{\mu}_{m'} - \hat{\mu}(k)\| &= \min_{1 \leq k \leq \bar{k}} \|\hat{\mu}_{m'} - \hat{\mu}_{m_k}\| \\ &\geq \min_{1 \leq k \leq \bar{k}} (-\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| - \|\mu(c_{m_k}) - \hat{\mu}_{m_k}\|) \\ &\geq \min_{1 \leq k \leq \bar{k}} \left(-2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| \right) \\ &= -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{1 \leq k \leq \bar{k}} \|\mu(c_{m'}) - \mu(c_{m_k})\| \\ &\geq -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\| \end{aligned}$$

where the last inequality follows from the fact that none of the existing centers shares the same true cluster index with arm m' .

The above upper bound and lower bound, together with the constraint on the accuracy of the empirical estimates, show that the \bar{k} -th center must come from $[M] \setminus \mathcal{A}_{\bar{k}}$. Therefore, Maximin Initialization succeeds in choosing K centers from K disjoint clusters.

Second, we prove that after the first step (Line 6) in the first iteration of Weighted K -means, \hat{c} is a correct partition, i.e., $\hat{c} \sim c$.

For any arm $m' \in [M]$, we can always find exactly one arm (denoted as $m_{\bar{k}}$) from m_1, m_2, \dots, m_K that shares the same true cluster index since m_1, m_2, \dots, m_K are selected from K disjoint clusters. The distance between $\hat{\mu}_{m'}$ and $\hat{\mu}(\bar{k})$ can be upper bounded as follows:

$$\begin{aligned} \|\hat{\mu}_{m'} - \hat{\mu}(\bar{k})\| &\leq \|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \hat{\mu}(\bar{k})\| \\ &= \|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m_{\bar{k}}}) - \hat{\mu}_{m_{\bar{k}}}\| \\ &\leq 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|. \end{aligned}$$

However, with respect to the remaining cluster centers, we have

$$\begin{aligned} \min_{k \in [K]: k \neq \bar{k}} \|\hat{\mu}_{m'} - \hat{\mu}(k)\| &\geq \min_{k \in [K]: k \neq \bar{k}} (-\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| - \|\mu(c_{m_k}) - \hat{\mu}(k)\|) \\ &= \min_{k \in [K]: k \neq \bar{k}} (-\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| - \|\mu(c_{m_k}) - \hat{\mu}_{m_k}\|) \\ &\geq \min_{k \in [K]: k \neq \bar{k}} \left(-2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \|\mu(c_{m'}) - \mu(c_{m_k})\| \right) \\ &= -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{k \in [K]: k \neq \bar{k}} \|\mu(c_{m_{\bar{k}}}) - \mu(c_{m_k})\| \\ &\geq -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\| \\ &> -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + 4 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| \end{aligned}$$

$$\begin{aligned}
 &= 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| \\
 &\geq \|\hat{\mu}_{m'} - \hat{\mu}(\bar{k})\|.
 \end{aligned}$$

Hence,

$$\hat{c}_{m'} = \arg \min_{k \in [K]} \|\hat{\mu}_{m'} - \hat{\mu}(k)\| = \bar{k}.$$

Since the arm m' is arbitrary, the above argument shows those arms that share the same true cluster index still share the same cluster index in \hat{c} . Besides, there are K disjoint non-empty clusters in \hat{c} . Therefore, \hat{c} is a correct partition.

Finally, we prove that the partition \hat{c} no longer changes after the first iteration of Weighted K -means, which we term as the Clustering is *stabilized*.

We need to show after the update (Line 7) of $\hat{\mu}(k)$, $\arg \min_{k \in [K]} \|\hat{\mu}_m - \hat{\mu}(k)\|$ still returns \hat{c}_m for any arm $m \in [M]$.

Since $\hat{c} \sim c$, there exists a permutation σ on $[K]$ such that $c = \sigma(\hat{c})$. For any cluster $\bar{k} \in [K]$,

$$\begin{aligned}
 \|\hat{\mu}(\bar{k}) - \mu(\sigma(\bar{k}))\| &= \left\| \frac{\sum_{m \in [M]} N_m \hat{\mu}_m \mathbb{1}\{\hat{c}_m = \bar{k}\}}{\sum_{m \in [M]} N_m \mathbb{1}\{\hat{c}_m = \bar{k}\}} - \mu(\sigma(\bar{k})) \right\| \\
 &= \left\| \frac{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\} (\hat{\mu}_m - \mu(\sigma(\bar{k})))}{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\}} \right\| \\
 &= \left\| \frac{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\} (\hat{\mu}_m - \mu(c_m))}{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\}} \right\| \\
 &\leq \frac{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\} \|\hat{\mu}_m - \mu(c_m)\|}{\sum_{m \in [M]} N_m \mathbb{1}\{c_m = \sigma(\bar{k})\}} \\
 &\leq \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|. \tag{24}
 \end{aligned}$$

Then for any arm $m' \in [M]$, the distance between $\hat{\mu}_{m'}$ and $\hat{\mu}(\hat{c}_{m'})$ can be upper bounded as follows

$$\begin{aligned}
 \|\hat{\mu}_{m'} - \hat{\mu}(\hat{c}_{m'})\| &\leq \|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \hat{\mu}(\hat{c}_{m'})\| \\
 &= \|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(\sigma(\hat{c}_{m'})) - \hat{\mu}(\hat{c}_{m'})\| \\
 &\leq 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|
 \end{aligned}$$

while the minimal distance between $\hat{\mu}_{m'}$ and other centers can be lower bounded as follows

$$\begin{aligned}
 &\min_{k \in [K]: k \neq \hat{c}_{m'}} \|\hat{\mu}_{m'} - \hat{\mu}(k)\| \\
 &\geq \min_{k \in [K]: k \neq \hat{c}_{m'}} (-\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(c_{m'}) - \mu(\sigma(k))\| - \|\mu(\sigma(k)) - \hat{\mu}(k)\|) \\
 &= \min_{k \in [K]: k \neq \hat{c}_{m'}} (-\|\hat{\mu}_{m'} - \mu(c_{m'})\| + \|\mu(\sigma(\hat{c}_{m'})) - \mu(\sigma(k))\| - \|\mu(\sigma(k)) - \hat{\mu}(k)\|) \\
 &\geq \min_{k \in [K]: k \neq \hat{c}_{m'}} \left(-2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \|\mu(\sigma(\hat{c}_{m'})) - \mu(\sigma(k))\| \right)
 \end{aligned}$$

$$\begin{aligned}
 &= -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{k \in [K]: k \neq \hat{c}_{m'}} \|\mu(\sigma(\hat{c}_{m'})) - \mu(\sigma(k))\| \\
 &\geq -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\| \\
 &> -2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| + 4 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| \\
 &= 2 \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\| \\
 &\geq \|\hat{\mu}_{m'} - \hat{\mu}(\hat{c}_{m'})\|.
 \end{aligned}$$

Therefore, $\arg \min_{k \in [K]} \|\hat{\mu}_{m'} - \hat{\mu}(k)\|$ remains equal to $\hat{c}_{m'}$. Since the arm m' is arbitrary, the partition is proved to be stabilized after the first iteration of Weighted K -means.

Moreover, $\hat{\mu}(k)$ for all $k \in [K]$ will stay the same once the partition stabilizes. Hence, by Equation (24),

$$\max_{k \in [K]} \|\hat{\mu}(k) - \mu(\sigma(k))\| \leq \max_{m \in [M]} \|\hat{\mu}_m - \mu(c_m)\|.$$

Now the proof of Proposition 10 is completed. \blacksquare

D.2 Proof of Proposition 12

Proof Due to the forced exploration in Algorithm 1 and the strong law of large numbers, for all $m \in [M]$, $\hat{\mu}_m(t)$ converges almost surely to $\mu(c_m)$, i.e., $\hat{\mu}_m(t) \xrightarrow{\text{a.s.}} \mu(c_m)$, as t tends to infinity.

In the following, we condition on the event

$$\mathcal{E} = \left\{ \lim_{t \rightarrow \infty} \hat{\mu}_m(t) = \mu(c_m) \text{ for all } m \in [M] \right\}$$

which has probability 1.

Note that in finite-dimensional spaces, pointwise convergence and convergence in Euclidean norm are equivalent. Therefore, by Proposition 10, we know that for sufficiently large t , Algorithm 2 will output a correct partition c^t such that $c^t = \sigma^t(c)$ for some permutation σ^t on $[K]$; moreover, $\mu^t(\sigma^t(k)) \rightarrow \mu(k)$ for all $k \in [K]$. Therefore, $\text{Alt}(c^t) = \text{Alt}(c)$ for sufficiently large t , and $\sigma^t(\mathcal{U}^t) \rightarrow \mathcal{U}$. Hence, for sufficiently large t ,

$$\begin{aligned}
 \lambda^*(t) &= \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c^t)} \sum_{m=1}^M \lambda_m \|\mu^t(c_m^t) - \mu'(c_m')\|^2 \\
 &= \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu^t(\sigma^t(c_m)) - \mu'(c_m')\|^2.
 \end{aligned}$$

By Proposition 9,

$$\begin{aligned}
 \lambda^*(t) &\rightarrow \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c_m')\|^2 \\
 &= \lambda^*.
 \end{aligned}$$

Consequently, $\lambda^*(t)$ converges pointwisely to λ^* . That is to say, for any $\epsilon > 0$, there exists $t_0(\epsilon)$ such that

$$\sup_{t > t_0(\epsilon)} \max_{m \in [M]} |\lambda_m^*(t) - \lambda_m^*| \leq \frac{\epsilon}{3(M-1)}.$$

By Lemma 22, there further exists $t_1(\epsilon) \geq t_0(\epsilon)$ such that

$$\sup_{t > t_1(\epsilon)} \max_{m \in [M]} \left| \frac{N_m(t)}{t} - \lambda_m^* \right| \leq \epsilon.$$

This is also equivalent to

$$\lim_{t \rightarrow \infty} \frac{N_m(t)}{t} = \lambda_m^* \text{ for all } m \in [M]$$

as desired. ■

D.3 Proof of Proposition 15

Proof Recall the definitions of $Z(t)$, $Z_1(t)$, and $Z_2(t)$ from Section 5.3 or Algorithm 1.

We can then bound the error probability as follows:

$$\begin{aligned} & \Pr(\tau_\delta < \infty, c^{\text{out}} \not\sim c) \\ & \leq \Pr(\exists t \in \mathbb{N} : c^{t-1} \not\sim c, Z(t) \geq \beta(\delta, t)) \\ & = \Pr\left(\exists t \in \mathbb{N} : c^{t-1} \not\sim c, \frac{1}{2} \max_{\alpha \geq 0} \left(-\alpha Z_1(t) + \frac{\alpha}{\alpha+1} Z_2(t)\right) \geq \beta(\delta, t)\right) \end{aligned} \quad (25)$$

$$\leq \Pr\left(\exists t \in \mathbb{N} : \frac{1}{2} \max_{\alpha \geq 0} \left(-\alpha Z_1(t) + \frac{\alpha}{\alpha+1} \sum_{m=1}^M N_m(t) \|\mu^{t-1}(c_m^{t-1}) - \mu(c_m)\|^2\right) \geq \beta(\delta, t)\right) \quad (26)$$

$$\leq \Pr\left(\exists t \in \mathbb{N} : \frac{1}{2} \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu(c_m)\|^2 \geq \beta(\delta, t)\right) \quad (27)$$

$$\leq \delta. \quad (28)$$

Line (25) follows from Lemma 23.

Line (26) follows from the fact that if $c^{t-1} \not\sim c$, then $(c, \mathcal{U}) \in \text{Alt}(c^{t-1})$ and hence

$$Z_2(t) \leq \sum_{m=1}^M N_m(t) \|\mu^{t-1}(c_m^{t-1}) - \mu(c_m)\|^2.$$

Line (27) follows from Lemma 24 and thus

$$\begin{aligned}
& \max_{\alpha \geq 0} \left(-\alpha Z_1(t) + \frac{\alpha}{\alpha + 1} \sum_{m=1}^M N_m(t) \|\mu^{t-1}(c_m^{t-1}) - \mu(c_m)\|^2 \right) \\
&= \max_{\alpha \geq 0} \left(\sum_{m=1}^M N_m(t) \left(-\alpha \|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2 + \frac{\alpha}{\alpha + 1} \|\mu^{t-1}(c_m^{t-1}) - \mu(c_m)\|^2 \right) \right) \\
&\leq \max_{\alpha \geq 0} \left(\sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu(c_m)\|^2 \right) \\
&= \sum_{m=1}^M N_m(t) \|\hat{\mu}_m(t) - \mu(c_m)\|^2.
\end{aligned}$$

Line (28) follows from Lemma 25. ■

D.4 Proof of Proposition 18

Proof Similarly to the proof of Proposition 12, in the following, we condition on the event

$$\mathcal{E} = \left\{ \lim_{t \rightarrow \infty} \hat{\mu}_m(t) = \mu(c_m) \text{ for all } m \in [M] \right\}$$

which has probability 1.

Note that in finite-dimensional spaces, pointwise convergence and convergence in Euclidean norm are equivalent. By Proposition 10, there exists $t_1 > 0$ such that for all $t \geq t_1$, Algorithm 2 will output a correct partition $c^t = \sigma^t(c)$ for some permutation σ^t on $[K]$; moreover, $\mu^t(\sigma^t(k)) \rightarrow \mu(k)$ for all $k \in [K]$. Therefore, for all $m \in [M]$, $\mu^{t-1}(c_m^{t-1}) \rightarrow \mu(c_m)$.

Accordingly, on the event \mathcal{E} , as t tends to infinity, $\|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2 \rightarrow 0$ for all $m \in [M]$. Since $N_m(t)/t$ is uniformly bounded by 1 for all $m \in [M]$, we have

$$\frac{Z_1(t)}{t} = \sum_{m=1}^M \frac{N_m(t)}{t} \|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2 \rightarrow 0.$$

Now we consider $Z_2(t)/t$. By Proposition 12 and its proof, conditioned on the event \mathcal{E} , $N_m(t)/t \rightarrow \lambda_m^*$ for all $m \in [M]$.

When $t > t_1$, c^{t-1} is always correct (i.e., $c^{t-1} \sim c$) and hence $\text{Alt}(c^{t-1}) = \text{Alt}(c)$. Thus, for $t > t_1$,

$$\begin{aligned}
\frac{Z_2(t)}{t} &= \min_{(c', \mathcal{U}') \in \text{Alt}(c^{t-1})} \sum_{m=1}^M \frac{N_m(t)}{t} \|\mu^{t-1}(c_m^{t-1}) - \mu'(c'_m)\|^2 \\
&= \min_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \frac{N_m(t)}{t} \|\mu^{t-1}(\sigma^{t-1}(c_m)) - \mu'(c'_m)\|^2.
\end{aligned}$$

By Proposition 6, as t tends to infinity,

$$\begin{aligned} \frac{Z_2(t)}{t} &\rightarrow \min_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m^* \|\mu(c_m) - \mu'(c'_m)\|^2 \\ &= 2D^*(c, \mathcal{U})^{-1}. \end{aligned}$$

Consequently,

$$\frac{Z(t)}{t} = \frac{1}{2} \left(\left(-\sqrt{\frac{Z_1(t)}{t}} + \sqrt{\frac{Z_2(t)}{t}} \right)_+ \right)^2 \rightarrow D^*(c, \mathcal{U})^{-1}.$$

So for any $0 < \epsilon < 1$, there exists $t_2 > t_1$ such that for all $t \geq t_2$,

$$\frac{Z(t)}{t} \geq (1 - \epsilon)D^*(c, \mathcal{U})^{-1}.$$

Now consider $\beta(\delta, t)$. Since $N_m(t) \leq t$ for $m \in [M]$, it holds that there exists $t_3 > 0$ such that for all $t \geq t_3$,

$$\sum_{m=1}^M 2d \log(4 + \log(N_m(t))) \leq \log(t)$$

which implies that

$$\beta(\delta, t) \leq \log(t) + Md \cdot \Psi \left(\frac{\log(1/\delta)}{Md} \right).$$

Altogether, we have

$$\begin{aligned} \tau_\delta &= \inf\{t \in \mathbb{N} : Z(t) \geq \beta(\delta, t)\} \\ &\leq \max\{t_2, t_3\} \vee \inf \left\{ t \in \mathbb{N} : t(1 - \epsilon)D^*(c, \mathcal{U})^{-1} \geq \log(t) + Md \cdot \Psi \left(\frac{\log(1/\delta)}{Md} \right) \right\} \end{aligned}$$

which shows τ_δ is finite conditioned on \mathcal{E} . Since $\Pr(\mathcal{E}) = 1$, we have

$$\Pr(\tau_\delta < \infty) = 1.$$

For ease of notation, let $a := (1 - \epsilon)D^*(c, \mathcal{U})^{-1}$ and $b := Md \cdot \Psi \left(\frac{\log(1/\delta)}{Md} \right)$. For sufficiently small δ , since $b \geq \log(1/\delta)$, $b + \log \left(\frac{1}{a} \right) > 0$. Thus, by Lemma 26,

$$\tau_\delta \leq \max \left\{ t_2, t_3, \frac{1}{a} \left(b + \log \left(\frac{e}{a} \right) + \log \left(b + \log \left(\frac{1}{a} \right) \right) \right) \right\}.$$

Note that t_2 and t_3 do not depend on δ and $\Psi(x) = x + \log(x) + o(\log(x))$ as $x \rightarrow \infty$. Thus,

$$\begin{aligned} \limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\log(1/\delta)} &\leq \limsup_{\delta \rightarrow 0} \frac{1}{a \log(1/\delta)} \left(b + \log \left(\frac{e}{a} \right) + \log \left(b + \log \left(\frac{1}{a} \right) \right) \right) \\ &= \frac{1}{a} \\ &= (1 - \epsilon)^{-1} D^*(c, \mathcal{U}). \end{aligned}$$

Since the above inequality holds for any $0 < \epsilon < 1$, by letting $\epsilon \rightarrow 0^+$,

$$\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\log(1/\delta)} \leq D^*(c, \mathcal{U}).$$

■

D.5 Proof of Theorem 19

Proof Let $0 < \epsilon < 1$.

First, we consider the sampling rule. According to Proposition 10 and Proposition 9, there exists a function

$$\xi : (0, 1) \rightarrow \left(0, \frac{1}{4} \min_{k, k' \in [K]: k \neq k'} \|\mu(k) - \mu(k')\|\right)$$

such that $\lim_{\epsilon \rightarrow 0^+} \xi(\epsilon) = 0$ and for any $t > 0$, if

$$\max_{m \in [M]} \|\hat{\mu}_m(t) - \mu(c_m)\| \leq \xi(\epsilon),$$

then:

- (i) Algorithm 2 outputs a correct partition c^t such that $c^t = \sigma^t(c)$ for some permutation σ^t on $[K]$;
- (ii) \mathcal{U}^t satisfies that

$$\max_{k \in [K]} \|\mu^t(\sigma^t(k)) - \mu(k)\| \leq \xi(\epsilon);$$

- (iii) $\lambda^*(t)$ satisfies that

$$\max_{m \in [M]} |\lambda_m^*(t) - \lambda_m^*| \leq \epsilon.$$

Point (iii) is, in fact, a consequence of the continuity of $\lambda^*(t)$ (see Proposition 9):

$$\begin{aligned} \lim_{\sigma^t(\mathcal{U}^t) \rightarrow \mathcal{U}} \lambda^*(t) &= \lim_{\sigma^t(\mathcal{U}^t) \rightarrow \mathcal{U}} \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c^t)} \sum_{m=1}^M \lambda_m \|\mu^t(c_m^t) - \mu'(c'_m)\|^2 \\ &= \lim_{\sigma^t(\mathcal{U}^t) \rightarrow \mathcal{U}} \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu^t(\sigma^t(c_m)) - \mu'(c'_m)\|^2 \\ &= \arg \max_{\lambda \in \mathcal{P}_M} \inf_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m \|\mu(c_m) - \mu'(c'_m)\|^2 \\ &= \lambda^*. \end{aligned}$$

In addition, let $T \geq \lceil M/\epsilon^3 \rceil \in \mathbb{N}$ and define the event

$$\mathcal{E}_T(\epsilon) = \bigcap_{t=\lceil T\epsilon^3 \rceil}^{\infty} \left\{ \max_{m \in [M]} \|\hat{\mu}_m(t) - \mu(c_m)\| \leq \xi(\epsilon) \right\}.$$

Conditioned on the event $\mathcal{E}_T(\epsilon)$, by Lemma 22 and the definition of $\xi(\epsilon)$, for all $t \geq T$,

$$\max_{m \in [M]} \left| \frac{N_m(t)}{t} - \lambda_m^* \right| \leq 3(M-1)\epsilon.$$

Now we introduce $D_\epsilon^*(c, \mathcal{U})^{-1}$, which is an ϵ -approximation of $D^*(c, \mathcal{U})^{-1}$, and defined as

$$D_\epsilon^*(c, \mathcal{U})^{-1} := \frac{1}{2} \inf_{\hat{\lambda}, \hat{\mathcal{U}}} \min_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \hat{\lambda}_m \|\hat{\mu}(c_m) - \mu'(c'_m)\|^2$$

where the infimum is over all $\hat{\lambda}$ and $\hat{\mathcal{U}}$ such that

$$\begin{cases} \max_{m \in [M]} |\hat{\lambda}_m - \lambda_m^*| \leq 3(M-1)\epsilon, \\ \max_{k \in [K]} \|\hat{\mu}(k) - \mu(k)\| \leq \xi(\epsilon). \end{cases}$$

By the continuity of $\epsilon \mapsto D_\epsilon^*(c, \mathcal{U})^{-1}$ at 0 (which is a consequence of Proposition 6),

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} D_\epsilon^*(c, \mathcal{U})^{-1} &= \frac{1}{2} \min_{(c', \mathcal{U}') \in \text{Alt}(c)} \sum_{m=1}^M \lambda_m^* \|\mu(c_m) - \mu'(c'_m)\|^2 \\ &= D^*(c, \mathcal{U})^{-1}. \end{aligned}$$

Therefore, conditioned on the event $\mathcal{E}_T(\epsilon)$, for all $t > T$, $Z_2(t)/t \geq 2D_\epsilon^*(c, \mathcal{U})^{-1}$. Concerning $Z_1(t)$, conditioned on the event $\mathcal{E}_T(\epsilon)$, for all $t > T$, using the inequality in Lemma 24 (with $\alpha = 1$),

$$\begin{aligned} \frac{Z_1(t)}{t} &= \sum_{m=1}^M \frac{N_m(t)}{t} \|\hat{\mu}_m(t) - \mu^{t-1}(c_m^{t-1})\|^2 \\ &\leq \sum_{m=1}^M \frac{2N_m(t)}{t} (\|\hat{\mu}_m(t) - \mu(c_m)\|^2 + \|\mu(c_m) - \mu^{t-1}(c_m^{t-1})\|^2) \\ &= \sum_{m=1}^M \frac{2N_m(t)}{t} (\|\hat{\mu}_m(t) - \mu(c_m)\|^2 + \|\mu(c_m) - \mu^{t-1}(\sigma^{t-1}(c_m))\|^2) \\ &\leq \sum_{m=1}^M \frac{2N_m(t)}{t} \left(\max_{m \in [M]} \|\hat{\mu}_m(t) - \mu(c_m)\| + \max_{k \in [K]} \|\mu^{t-1}(\sigma^{t-1}(k)) - \mu(k)\| \right) \\ &\leq 4\xi(\epsilon). \end{aligned}$$

Altogether, conditioned on $\mathcal{E}_T(\epsilon)$, for all $t > T$,

$$\begin{aligned} \frac{Z(t)}{t} &= \frac{1}{2} \left(\left(-\sqrt{\frac{Z_1(t)}{t}} + \sqrt{\frac{Z_2(t)}{t}} \right)_+ \right)^2 \\ &\geq \left(\left(\sqrt{D_\epsilon^*(c, \mathcal{U})^{-1}} - \sqrt{2\xi(\epsilon)} \right)_+ \right)^2. \end{aligned}$$

Then consider $\beta(\delta, t)$. Since $N_m(t) \leq t$ for $m \in [M]$, it holds that there exists $t_2 > 0$ such that for all $t \geq t_2$,

$$\sum_{m=1}^M 2d \log(4 + \log(N_m(t))) \leq \log(t)$$

which implies that

$$\beta(\delta, t) \leq \log(t) + Md \cdot \Psi\left(\frac{\log(1/\delta)}{Md}\right).$$

For ease of notation, let $a := \left(\left(\sqrt{D_\epsilon^*(c, \mathcal{U})^{-1}} - \sqrt{2\xi(\epsilon)}\right)_+\right)^2$ and $b := Md \cdot \Psi\left(\frac{\log(1/\delta)}{Md}\right)$. Consequently, conditioned on $\mathcal{E}_T(\epsilon)$,

$$\begin{aligned} \tau_\delta &= \inf\{t \in \mathbb{N} : Z(t) \geq \beta(\delta, t)\} \\ &\leq \max\{T + 1, t_2\} \vee \inf\{t \in \mathbb{N} : at \geq \log(t) + b\}. \end{aligned}$$

Let $t_3 := \max\{\lceil M/\epsilon^3 \rceil + 1, t_2\} \vee \inf\{t \in \mathbb{N} : at \geq \log(t) + b\}$. Then for all $T + 1 \geq t_3$, $\mathcal{E}_T(\epsilon) \subseteq \{\tau_\delta \leq T + 1\}$ and hence $\Pr(\tau_\delta > T + 1) \leq \Pr(\mathcal{E}_T^c(\epsilon))$. Therefore, $\mathbb{E}[\tau_\delta]$ can be upper bounded as follows:

$$\begin{aligned} \mathbb{E}[\tau_\delta] &= \sum_{t=0}^{\infty} \Pr(\tau_\delta > t) \\ &= \sum_{t=0}^{t_3-1} \Pr(\tau_\delta > t) + \sum_{t=t_3}^{\infty} \Pr(\tau_\delta > t) \\ &\leq t_3 + \sum_{T=t_3-1}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon)) \\ &\leq t_3 + \sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon)). \end{aligned}$$

Since $\lim_{\epsilon \rightarrow 0^+} a = D^*(c, \mathcal{U})^{-1}$, $a > 0$ for sufficiently small ϵ . Together with $b \geq \log(1/\delta)$, we have $b + \log\left(\frac{1}{a}\right) > 0$ for sufficiently small δ . Thus, by Lemma 26,

$$\mathbb{E}[\tau_\delta] \leq \max\left\{\left\lceil \frac{M}{\epsilon^3} \right\rceil + 1, t_2, \frac{1}{a} \left(b + \log\left(\frac{e}{a}\right) + \log\left(b + \log\left(\frac{1}{a}\right)\right)\right)\right\} + \sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon)).$$

For the moment, let us assume that the final term $\sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon))$ is finite. Note that $\lceil M/\epsilon^3 \rceil$, t_2 and $\sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon))$ do not depend on δ and $\Psi(x) = x + \log(x) + o(\log(x))$ as $x \rightarrow \infty$. Thus,

$$\begin{aligned} \limsup_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)} &\leq \limsup_{\delta \rightarrow 0} \frac{1}{a \log(1/\delta)} \left(b + \log\left(\frac{e}{a}\right) + \log\left(b + \log\left(\frac{1}{a}\right)\right)\right) \\ &= \frac{1}{a} \\ &= \left(\left(\sqrt{D_\epsilon^*(c, \mathcal{U})^{-1}} - \sqrt{\xi(\epsilon)}\right)_+\right)^{-2}. \end{aligned}$$

Since the above inequality holds for any $\epsilon > 0$ such that $a > 0$, by letting $\epsilon \rightarrow 0^+$,

$$\limsup_{\delta \rightarrow 0} \frac{\tau_\delta}{\log(1/\delta)} \leq D^*(c, \mathcal{U}).$$

It remains to show that $\sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon))$ is finite. Since

$$\mathcal{E}_T^c(\epsilon) = \bigcup_{t=\lceil T\epsilon^3 \rceil}^{\infty} \left\{ \max_{m \in [M]} \|\hat{\mu}_m(t) - \mu(c_m)\| > \xi(\epsilon) \right\},$$

by a union bound, we have

$$\begin{aligned} \Pr(\mathcal{E}_T^c(\epsilon)) &\leq \sum_{t=\lceil T\epsilon^3 \rceil}^{\infty} \sum_{m \in [M]} \Pr(\|\hat{\mu}_m(t) - \mu(c_m)\| > \xi(\epsilon)) \\ &\leq \sum_{t=\lceil T\epsilon^3 \rceil}^{\infty} \sum_{m \in [M]} \Pr\left(\exists i \in [d] : |[\hat{\mu}_m(t) - \mu(c_m)]_i| > \frac{\xi(\epsilon)}{\sqrt{d}}\right) \\ &\leq \sum_{t=\lceil T\epsilon^3 \rceil}^{\infty} \sum_{m \in [M]} \sum_{i \in [d]} \Pr\left(|[\hat{\mu}_m(t) - \mu(c_m)]_i| > \frac{\xi(\epsilon)}{\sqrt{d}}\right), \end{aligned}$$

where we use $[x]_i$ to denote the i^{th} component of the vector $x \in \mathbb{R}^d$.

Considering sufficiently large T , for any $t \geq \lceil T\epsilon^3 \rceil$, $N_m(t) \geq \sqrt{t}/2$ for all $m \in [M]$. Then, by Hoeffding's inequality for sub-Gaussian random variables,

$$\begin{aligned} \Pr(\mathcal{E}_T^c(\epsilon)) &\leq 2Md \sum_{t=\lceil T\epsilon^3 \rceil}^{\infty} \exp\left(-\frac{\xi(\epsilon)^2 \sqrt{t}}{d}\right) \\ &\leq 2Md \int_{x=\lceil T\epsilon^3 \rceil - 1}^{\infty} \exp\left(-\frac{\xi(\epsilon)^2 \sqrt{x}}{d}\right) dx \\ &= \frac{4Md^3}{\xi(\epsilon)^4} \left(\frac{\xi(\epsilon)^2 \sqrt{\lceil T\epsilon^3 \rceil - 1}}{d} + 1\right) \exp\left(-\frac{\xi(\epsilon)^2 \sqrt{\lceil T\epsilon^3 \rceil - 1}}{d}\right) \\ &\leq \frac{4Md^3}{\xi(\epsilon)^4} \left(\frac{\xi(\epsilon)^2 \sqrt{T\epsilon^3}}{d} + 1\right) \exp\left(-\frac{\xi(\epsilon)^2 \sqrt{T\epsilon^3}}{2d}\right). \end{aligned}$$

Since for any constant $c > 0$, the infinite sum $\sum_{n=1}^{\infty} \sqrt{n} \exp(-c\sqrt{n})$ is convergent, $\sum_{T=\lceil M/\epsilon^3 \rceil}^{\infty} \Pr(\mathcal{E}_T^c(\epsilon))$ is also convergent.

This completes the proof of Theorem 19. ■

Appendix E. Additional Numerical Results on a Higher-Dimensional Dataset

In this appendix, we extend our numerical experiments to the MNIST dataset (LeCun et al., 1998), which comprises data points of higher dimensions. Specifically, we conduct

d	BOC	UNIFORM
10	2932.0 \pm 52.4	6991.8 \pm 120.8
50	8782.6 \pm 110.8	16418.7 \pm 173.8
100	16663.2 \pm 162.8	31270.8 \pm 273.5
200	32840.5 \pm 138.8	61543.7 \pm 275.3
400	65414.2 \pm 209.6	122508.3 \pm 369.9

Table 3: The averaged empirical sample complexities of BOC and UNIFORM for $\delta = 0.01$ on the MNIST dataset.

our online clustering task as follows. First, we randomly sample 500 images of digits from the dataset, vectorize them, and, akin to Sheikh et al. (2020), we employ Principal Component Analysis (PCA) to project the data onto different number of dimensions. Subsequently, we preprocess the dataset consistently with our approach to the Iris and Yeast datasets (see Section 6). This yields the number of clusters $K = 10$ and the number of arms $M = 500$, with the problem dimension d set to be in $\{10, 50, 100, 200, 400\}$. Note that the original dimensionality of the data points in the MNIST dataset is $784 = 28 \times 28$.

The sample complexities for confidence level $\delta = 0.01$ are presented in Table 3. It is evident from the table that BOC outperforms the non-adaptive baseline method, UNIFORM, by a factor of ≈ 2 in terms of their sample complexities, underscoring the advantage of employing an adaptive sampling approach, even on this dataset with large dimensionality.

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24, 2011.
- Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18:6446–6531, 2017.
- Shipra Agrawal and Navin Goyal. Analysis of Thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory (COLT)*, pages 39–1. JMLR Workshop and Conference Proceedings, 2012.
- Daniel Aloise, Amit Deshpande, Pierre Hansen, and Preyas Popat. NP-hardness of Euclidean sum-of-squares clustering. *Machine Learning*, 75(2):245–248, 2009.
- D. Arthur and S. Vassilvitskii. K-means++: The advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1027–1035, 2007.
- Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *Conference on Learning Theory (COLT)*, pages 41–53, 2010.

- Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- C. Berge. *Topological Spaces: Including a Treatment of Multi-valued Functions, Vector Spaces, and Convexity*. Oliver & Boyd, 1963.
- Djallel Bouneffouf, Srinivasan Parthasarathy, Horst Samulowitz, and Martin Wistuba. Optimal exploitation of clustering and history information in multi-armed bandit. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 2016–2022, 2019.
- Sébastien Bubeck and Nicolò Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1): 1–122, 2012.
- Emil Carlsson, Devdatt Dubhashi, and Fredrik D. Johansson. Thompson sampling for bandits with clustered arms. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, pages 2212–2218, 2021.
- Raymond B. Cattell. The description of personality: Basic traits resolved into clusters. *The Journal of Abnormal and Social Psychology*, 38(4):476, 1943.
- M Emre Celebi, Hassan A Kingravi, and Patricio A Vela. A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Systems with Applications*, 40(1):200–210, 2013.
- Anil Chaturvedi, J Douglas Carroll, Paul E Green, and John A Rotondo. A feature-based approach to market segmentation via overlapping k-centroids clustering. *Journal of Marketing Research*, 34(3):370–377, 1997.
- Herman Chernoff. Sequential design of experiments. *The Annals of Mathematical Statistics*, 30(3):755–770, 1959.
- Anna Choromanska and Claire Monteleoni. Online clustering with experts. In *International Conference on Artificial Intelligence and Statistics*, pages 227–235. PMLR, 2012.
- Vincent Cohen-Addad, Benjamin Guedj, Varun Kanade, and Guy Rom. Online k-means clustering. In *International Conference on Artificial Intelligence and Statistics*, pages 1126–1134. PMLR, 2021.
- Harold Edson Driver and Alfred Louis Kroeber. *Quantitative Expression of Cultural Relationships*, volume 31. Berkeley: University of California Press, 1932.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7(6), 2006.

- Yifan Feng, Rene Caldentey, and Christopher Thomas Ryan. Robust learning of consumer preferences. *Operations Research*, 2021.
- Edward W. Forgy. Cluster analysis of multivariate data: Efficiency versus interpretability of classifications. *Biometrics*, 21:768–769, 1965.
- Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pages 998–1027. PMLR, 2016.
- Manlio Gaudioso, Giovanni Giallombardo, and Giovanna Miglionico. An incremental method for solving convex finite min-max problems. *Mathematics of Operations Research*, 31(1):173–187, 2006.
- Claudio Gentile, Shuai Li, and Giovanni Zappella. Online clustering of bandits. In *International Conference on Machine Learning*, pages 757–765. PMLR, 2014.
- Cristina Gigola and Susana Gomez. A regularization method for solving the finite convex min-max problem. *SIAM Journal on Numerical Analysis*, 27(6):1621–1634, 1990.
- Teofilo F. Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.
- Ronald L. Graham, Donald E. Knuth, Oren Patashnik, and Stanley Liu. Concrete mathematics: A foundation for computer science. *Computers in Physics*, 3(5):106–107, 1989.
- Jeffrey W. Herrmann. A genetic algorithm for minimax optimization problems. In *Proceedings of the 1999 Congress on Evolutionary Computation-CEC99 (Cat. No. 99TH8406)*, volume 2, pages 1099–1103. IEEE, 1999.
- Mary Inaba, Naoki Katoh, and Hiroshi Imai. Applications of weighted voronoi diagrams and randomization to variance-based k-clustering. In *Proceedings of the Tenth Annual Symposium on Computational Geometry*, pages 332–339, 1994.
- Anil K. Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010.
- Anil K. Jain, M Narasimha Murty, and Patrick J Flynn. Data clustering: A review. *ACM Computing Surveys (CSUR)*, 31(3):264–323, 1999.
- Yassir Jedra and Alexandre Proutiere. Optimal best-arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 33:10007–10017, 2020.
- Kwang-Sung Jun, Kevin Jamieson, Robert Nowak, and Xiaojin Zhu. Top arm identification in multi-armed bandits with batch arm pulls. In *International Conference on Artificial Intelligence and Statistics*, pages 139–148. PMLR, 2016.
- Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pages 1238–1246. PMLR, 2013.

- P. N. Karthik and Rajesh Sundaresan. Learning to detect an odd Markov arm. *IEEE Transactions on Information Theory*, 66(7):4324–4348, 2020.
- P. N. Karthik and Rajesh Sundaresan. Detecting an odd restless Markov arm with a trembling hand. *IEEE Transactions on Information Theory*, 67(8):5230–5258, 2021.
- Emilie Kaufmann and Wouter M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- Azadeh Khaleghi, Daniil Ryabko, Jérémie Mary, and Philippe Preux. Online clustering of processes. In *International Conference on Artificial Intelligence and Statistics*, pages 601–609. PMLR, 2012.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- Tor Lattimore and Csaba Szepesvári. The end of optimism? An asymptotic analysis of finite-armed linear bandits. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 728–737. PMLR, 2017.
- Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Shuai Li, Alexandros Karatzoglou, and Claudio Gentile. Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 539–548, 2016.
- Yonglong Li and Vicent Y. F. Tan. Second-order asymptotics of sequential hypothesis testing. *IEEE Transactions on Information Theory*, 66(11):7222–7230, 2020.
- Edo Liberty, Ram Sriharsha, and Maxim Sviridenko. An algorithm for online k-means clustering. In *2016 Proceedings of the eighteenth workshop on algorithm engineering and experiments (ALENEX)*, pages 81–89. SIAM, 2016.
- Stuart Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- John David MacCuish and Norah E. MacCuish. *Clustering in Bioinformatics and Drug Discovery*. CRC Press, 2010.
- James MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA, 1967.

- Meena Mahajan, Prajakta Nimbhorkar, and Kasturi Varadarajan. The planar k-means problem is NP-hard. *Theoretical Computer Science*, 442:13–21, 2012.
- Mikhail B. Maluyutov and Ivan I. Tsitovich. *Second-Order Optimal Sequential Tests*, volume 51, chapter 7, pages 201–213. Atkinson A., Bogacka B., Zhigljavsky A. (eds) Optimum Design, Springer, Boston, MA., 2001.
- Arya Mazumdar and Barna Saha. Clustering with noisy queries. *Advances in Neural Information Processing Systems*, 30, 2017.
- Trong T. Nguyen and Hady W. Lauw. Dynamic clustering of contextual multi-armed bandits. In *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management*, pages 1959–1962, 2014.
- Gayathri R. Prabhu, Srikrishna Bhashyam, Aditya Gopalan, and Rajesh Sundaresan. Sequential multi-hypothesis testing in multi-armed bandit problems: An approach for asymptotic optimality. *arXiv preprint arXiv:2007.12961*, 2020.
- Clémence Réda, Andrea Tirinzoni, and Rémy Degenne. Dealing with misspecification in fixed-confidence linear top-m identification. *Advances in Neural Information Processing Systems*, 34, 2021.
- Enrique H. Ruspini. A new approach to clustering. *Information and Control*, 15(1):22–32, 1969.
- Amit Saxena, Mukesh Prasad, Akshansh Gupta, Neha Bharill, Om Prakash Patel, Aruna Tiwari, Meng Joo Er, Weiping Ding, and Chin-Teng Lin. A review of clustering techniques and developments. *Neurocomputing*, 267:664–681, 2017.
- Ruksar Sheikh, Mayank Patel, and Amit Sinhal. *Recognizing MNIST Handwritten Data Set Using PCA and LDA*, pages 169–177. In book: International Conference on Artificial Intelligence: Advances and Applications 2019, Feb 2020.
- Rahul Singh, Fang Liu, Yin Sun, and Ness Shroff. Multi-armed bandits with dependent arms. *arXiv preprint arXiv:2010.09478*, 2020.
- Marta Soare, Alessandro Lazaric, and Rémi Munos. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems 27*, 2014.
- Ting Su and Jennifer G Dy. In search of deterministic methods for initializing k-means and Gaussian mixture clustering. *Intelligent Data Analysis*, 11(4):319–338, 2007.
- Rangarajan K. Sundaram. *A First Course in Optimization Theory*. Cambridge University Press, 1996.
- William R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3-4):285–294, 1933.
- Nidhin Koshy Vaidhiyan and Rajesh Sundaresan. Learning to detect an oddball target. *IEEE Transactions on Information Theory*, 64(2):831–852, 2017.

Zhenlin Wang and Jonathan Scarlett. Max-min grouped bandits. In *Proc. of the 36th AAAI Conference on Artificial Intelligence*. AAAI Press, 2022.

Junwen Yang and Vincent Y. F. Tan. Minimax optimal fixed-budget best arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 35:12253–12266, 2022.

Zixin Zhong, Wang Chi Cheung, and Vincent Y. F. Tan. Best arm identification for cascading bandits in the fixed confidence setting. In *International Conference on Machine Learning*, pages 11481–11491. PMLR, 2020.