

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 ONLINE LEARNING WITH RECENCY: ALGORITHMS FOR SLIDING-WINDOW STREAMING MULTI-ARMED BANDITS

Anonymous authors

Paper under double-blind review

ABSTRACT

Motivated by the recency effect in online learning, we study algorithms for single-pass *sliding-window streaming multi-armed bandits (MABs)* in this paper. In this setting, we are given n arms with unknown sub-Gaussian reward distributions and a parameter W . The arms arrive in a single-pass stream, and only the most recent W arms are considered valid. The algorithm is required to perform pure exploration and regret minimization with *limited memory*, **defined as the number of stored arms**. The model is a natural extension of the streaming multi-armed bandits model (without the sliding window) that has been extensively studied in recent years. We provide a comprehensive analysis of both the pure exploration and regret minimization problems with the model. For pure exploration, we prove that finding the best arm is hard with sublinear memory while finding an *approximate* best arm admits an efficient algorithm. For regret minimization, we explore a new notion of regret and give sharp memory-regret trade-offs for any single-pass algorithms. We complement our theoretical results with experiments, demonstrating the trade-offs between sample, regret, and memory.

1 INTRODUCTION

The stochastic multi-armed bandits (MABs) model is a fundamental model extensively studied in machine learning (ML) and theoretical computer science (TCS). In its most common form, we are given n arm with unknown sub-Gaussian reward distributions, and we can learn the instance by *sampling* from the arms. The most important problems in the model include *pure exploration*, where the goal is to identify the best or a near-optimal arm, and *regret minimization*, where the aim is to devise a sampling strategy that performs competitively against the best arm in hindsight. The multi-armed bandits model has found broad applications in experiment design and clinical trials (Robbins, 1952; Pallmann et al., 2018; Simchi-Levi & Wang, 2023), financial strategies (Shen et al., 2015; Trovò et al., 2018), information retrieval (Radlinski et al., 2008; Losada et al., 2017), algorithm design (Bouneffouf et al., 2017; Gullo et al., 2023), to name a few.

Classical algorithms for MABs often assume the entire set of n is stored in the memory for repeated access. However, this assumption can be unrealistic in modern online learning and large-scale applications, where arms may arrive sequentially in a stream, and the available memory is insufficient to store all of them. To address this challenge, the work of Liau et al. (2018); Assadi & Wang (2020) introduced the *streaming* multi-armed bandits model. In this model, the arms arrive one after another in a stream, and the algorithm would ideally maintain a memory substantially smaller than the total number of arms. The maximum number of arms maintained in the memory is defined as the *space complexity* of the algorithm. The streaming MABs model has attracted considerable attention since its introduction, and a flurry of work has established near-tight trade-offs for pure exploration (Assadi & Wang, 2020; Jin et al., 2021; Maiti et al., 2021; Assadi & Wang, 2022; 2024; Karpov & Wang, 2025) and regret minimization (Liau et al., 2018; Maiti et al., 2021; Agarwal et al., 2022; Wang, 2023; He et al., 2025) in various settings.

047 While most work on streaming MABs targets global objectives, such as identifying the best arm overall, many
 048 applications exhibit a recency effect, where recent arms matter more. For example, movie recommendation
 049 systems must adapt quickly to shifting trends. A related motivation comes from privacy constraints: regula-
 050 tions and policies often mandate data deletion after limited periods. GDPR requires data retention only for
 051 the “necessary” duration (GDPR, 2016), Apple retains user data for 6 months (Apple Inc., 2021), and Google
 052 limits anonymized advertising data to 9 months (Google LLC, 2025). Alas, streaming MABs algorithms
 053 usually do not take any recency effect into consideration. For instance, the pure exploration algorithms, e.g.,
 054 the ones in Assadi & Wang (2020); Jin et al. (2021); Maiti et al. (2021), may output an arm that arrives very
 055 early in the stream, which is far from being recent. Similarly, the regret minimization algorithms in Maiti
 056 et al. (2021); Wang (2023); He et al. (2025) may commit to an arm that is outside the pool of recent arms¹.
 057 As such, the following motivating open question could be asked: *could we design efficient streaming MABs*
 058 *algorithms that incorporate the recency effect?*

059 **Sliding-window streaming multi-armed bandits.** One of the most common models that capture the recency
 060 effect is the sliding-window streaming model (Datar et al., 2002; Datar & Motwani, 2016). In a typical
 061 sliding-window stream, a total of n data items (arms in the context of MABs) are arriving in a stream, and
 062 only the past W items are considered valid. The sliding-window streams have been extensively studied in
 063 various contexts, including frequency estimation (Datar et al., 2002; Braverman & Ostrovsky, 2007), graph
 064 algorithms (Crouch et al., 2013; Crouch & Stubbs, 2014; Zhang et al., 2024), clustering (Braverman et al.,
 065 2016; Borassi et al., 2020; Epasto et al., 2022; Woodruff et al., 2023; Cohen-Addad et al., 2025), among
 066 others (Tao & Papadimas, 2006; Zhang et al., 2016).

067 Inspired by the success of sliding-window streams on various problems, we define the natural notion of
 068 sliding-window streaming MABs to explore the recency effect. Here, we are given n arms arriving in a
 069 (single-pass) stream, and we are additionally given a window size W . When the t -th arm arrives, the arms
 070 with the arrival orders in $[t - W + 1, t]$ are considered the *valid* set of arms at this point. **We emphasize that**
 071 **throughout the paper, W and n are parameters given by the problem instance, and we cannot adjust these**
 072 **parameters.** The algorithm is allowed to store *any* arm (not limited to the window) regardless of whether
 073 the arm is valid². The central problems here are therefore the *pure exploration* and *regret minimization* in
 074 sliding-window streaming MABs.

075 1.1 OUR CONTRIBUTIONS

076 We give a comprehensive analysis of pure exploration and regret minimization algorithms for sliding-window
 077 streaming MABs in this paper. Our results can be summarized in Table 1.

078 **Pure explorations.** For pure explorations, we studied both *pure exploration*, where the goal is to return the
 079 *exact* best arm, and ε exploration, where the goal is to return an arm whose mean is ε -close to the best. In
 080 both notions, the best arm is defined as the arm with the highest mean reward in the *sliding window*. Our
 081 main conceptual message is that finding the *exact* best arm is hard unless using $\Omega(W)$ arms of memory space,
 082 but finding the *approximation* best arm is possible with sample and space efficiency.

083 **Result 1** (Informal of Theorems 1 and 2). *The following statements are true for exploration in sliding-*
 084 *window MABs.*

- 085 • Any algorithm that finds the *(exact)* best arm at any step with probability at least 99/100 in
 086 the sliding-window streaming multi-armed bandits requires $\Omega(W)$ arm memory, even with an
 087 unlimited number of arm pulls.

088 089 090 091 ¹This intuitively means the algorithm incurs large regret, although the definition of regret has more nuance in such
 092 cases. See Section 1.1 and Section 2 for details.

093 ²The arms outside the sliding window could still be useful in various subroutines, e.g., comparing the means.

094

- There exists an algorithm that finds a (*approximate*) ε -best arm with probability at least $1 - \delta$

095 at any steps with $O(\frac{1}{\varepsilon})$ arm memory and $O(\frac{n}{\varepsilon^2} \log \frac{W}{\delta})$ arm pulls.

096

097 By a standard probability boosting argument, the success probability of 99/100 **in the lower bound** generalizes
 098 to *any* probability of $1/2 + \Omega(1)$. On the other hand, our results demonstrate that we can identify an
 099 approximate best arm with arbitrary constant accuracy using only $O(\frac{1}{\varepsilon})$ memory. **For constant choice of ε**
 100 (which is usually the case), our algorithm achieves *constant memory* for exploration.

101 **Regret minimization.** For *regret minimization*, a significant challenge is how to *define* regret in the sliding-
 102 window model. The most natural definition is to define the regret as the cumulative gap between $\mu^*(t, W)$
 103 and the means of the pulled arms in each window. Here, $\mu^*(t, W)$ is the mean reward of the optimal arm in
 104 the window W at time t . However, such a definition has a fatal issue: since the algorithm controls the number
 105 of arm pulls before the window moves, the definition of the regret becomes a function of the algorithm, which
 106 means it cannot be well-defined.

107 To bypass the issue, we introduce the notion of *epoch-wise* regret such that the optimal reward sequences
 108 are *independent* of the arm pulls used by the algorithm. Our notion of regret minimization is to divide the
 109 total number of arms pulls T into *equal-sized epochs*. Among the $n - W + 1$ epochs (one for each window
 110 position), each epoch is allocated with $\frac{T}{n-W+1}$ arm pulls. Total regret is defined as cumulative regret across
 111 epochs, and the algorithm is required to pull arms a constrained number of times in each time window. A
 112 formal definition of our regret notion can be found in Definition 6.

113 We believe that the introduction of the regret notion is a significant contribution; otherwise, there is no
 114 obvious way to study regret minimization in sliding-window streaming MABs. Moreover, the epoch-wise
 115 regret definition captures many practical scenarios. For instance, in the case of movie recommendations, we
 116 treat the “sliding window” as time periods of, e.g., 1-2 months, and we aim to recommend the most relevant
 117 movies in each period. Our main conceptual finding for regret minimization is that a memory of $\Omega(W)$ arms
 118 is necessary to achieve $o(T)$ regret; furthermore, there is a sharp memory-regret transition around the $\Theta(W)$
 119 arm memory.

120 **Result 2** (Informal of Theorem 3). *Any algorithm that achieves $o(T/W^2)$ regret in the epoch-wise
 121 regret setting requires $\Omega(W)$ arm memory. Furthermore, there exist algorithms that given a stream of n
 122 arms and parameters T and W , with $O(W)$ memory achieve $O(\sqrt{W \cdot (n - W) \cdot T})$ regret.*

124 In the centralized setting, the tight bound for regret minimization is $O(\sqrt{nT})$, even with unlimited memory.
 125 Since $O(\sqrt{W(n - W)T}) = O(\sqrt{nT})$ when $|n - W|$ or W is small, this shows that our bound for *epoch-wise*
 126 *regret* setting is indeed tight in the worst case. We find the conceptual message quite interesting, and we
 127 believe it could serve as important guidelines for related applications. A variant of our regret setting is when
 128 the best arm does not expire with the movement of the sliding window. While the setting is less interesting,
 129 we do believe it has applications as well. A discussion of this setting can be found in ??.

131 **Our techniques.** We start with ε -exploration in sliding-window streaming MABs, in which there are two
 132 main technical challenges: the memory constraints and the expiration of arms. An efficient sliding-window
 133 algorithm would imply an efficient streaming algorithm (by setting $W = n$); as such, for any MABs technique
 134 to work in the sliding-window setting, there must be a streaming algorithm as well. For the majority of MABs
 135 techniques, efficient streaming algorithms either do not exist (e.g., elimination-based algorithms Even-Dar
 136 et al. (2006); Karnin et al. (2013)), or it is unclear how to design such algorithms (e.g., non-stationary bandits
 137 Whittle (1988) and mortal bandits Chakrabarti et al. (2008)). Therefore, we have to find technical ideas from
 138 existing streaming MABs algorithms (e.g. Assadi & Wang (2020); Jin et al. (2021); Maiti et al. (2021)).

139 Most of the algorithms in streaming MABs are either based on amortizing the sample complexity across the
 140 stream or using a bucket-based idea group arms based on the empirical means. The idea of amortization faces

Task	Space	Sample/Regret	Remark
Exact Exploration	$\Omega(W)$	Any	Lower Bound
Strong ε -exploration	$O(1/\varepsilon)$	$\Theta(\frac{n}{\varepsilon^2} \log n)$	Upper and Lower Bounds
Weak ε -exploration	$O(1/\varepsilon)$	$\Theta(\frac{n}{\varepsilon^2} \log W)$	Upper Bound
Regret minimization	$o(W)$	$\Omega(T/W^2)$	Lower Bound
	$\Omega(W)$	$O(\sqrt{W \cdot (n - W) \cdot T})$	Upper Bound

Table 1: Summary of the results for exploration and regret minimization

a barrier aimed at the expiration of arms. In particular, for the algorithms that amortize sample complexity in, e.g., Assadi & Wang (2020); Jin et al. (2021), the guarantees are only given with respect to the best arm, and the analysis falls apart if the best arm changes due to the sliding window movements. On the other hand, the grouping of empirical means based on the multiplicative of ε is naturally compatible with the sliding-window model. Here, we can simply discard the expired arms, and the invariant among the buckets helps maintain ε -best arms. This is the main idea for our ε -exploration algorithms.

Our lower bound for the exact pure exploration establishes a sharp dichotomy between the exact and approximate ε -exploration problems. Here, the main challenge is to find a distribution that forces the algorithm to store virtually all arms. Our idea is to use a distribution of arms with *decreasing mean rewards*. This distribution forces the “useless” arm when the window is at position t to become optimal when the window moves to $t + 1$. Although the distribution is not involved, it crucially uses the sliding-window property to separate from the streaming case, especially given that the latter admits an efficient algorithm with a single-arm memory (Assadi & Wang (2020)). Finally, our regret lower bound follows the same idea, although we need to extend the distribution to slightly more involved ones to ensure the algorithm cannot get “lucky” with the instance distribution.

Experiments. We conducted experiments for both pure exploration and regret minimization applications³. For pure exploration, we implemented the ε -best pure exploration algorithm, and for regret minimization, we used the $O(W)$ -memory algorithms outlined in Result 2. These are the first algorithms designed to work with multi-armed bandits (MABs) under a sliding-window setting.

In our pure exploration experiments, we tested configurations with $n \in \{1000, 2000, 5000\}$ and $n \in \{10, 20, 50\}$. The results indicate a relatively smooth trade-off between the quality of the returned arm and the memory used. The error exceeded 0.6 in all settings when we employed a memory size of $0.05W$; however, it dropped to below 0.3 with a memory size of $0.3W$. On the other hand, we can easily show that existing algorithms could result in 0.6 error (??), and our empirical results essentially mean that with $0.3W$ memory, the error could be reduced by 50%. For the regret minimization experiments, we tested configurations with $n \in \{500, 1000, 2000\}$ and $n \in \{10, 20, 50\}$, while setting the number of pulls for each epoch to $\frac{T}{n-W+1} = 1000$. The results revealed sharp changes in regret around the memory size W , confirming our theoretical predictions. The total regret decreased by more than 50% for most configurations when the memory size increased from $0.05W$ to W .

³Our code is available on anonymous Github: <https://anonymous.4open.science/r/sliding-window-MABs-CF74/>.

188 **2 PROBLEM DEFINITION AND PRELIMINARIES**
189

190 In this section, we give the formal definition of the problems we investigated and some standard technical
191 tools. We start with a formal definition of stochastic MABs.

192 **Definition 1** (Stochastic multi-armed bandits (MABs) model). In the stochastic multi-armed bandits model,
193 we have a collection of n arms $\{\text{arm}_i\}_{i=1}^n$, and each arm follows a distribution with mean $\mu_i \in [0, 1]$. Each
194 pull of arm_i returns a sample from the distribution with mean μ_i .

195 Note that by the central limit theorem, sampling from arbitrary distributions over $[0, 1]$ is essentially the
196 same as sampling from an arbitrary sub-Gaussian distribution (up to a scaling factor). The sliding-window
197 streaming MABs could therefore be defined as follows.

198 **Definition 2** (The sliding-window streaming MABs model.). In the sliding-window streaming MABs model,
199 we have a collection of n arms $\{\text{arm}_i\}_{i=1}^n$ arranged in order and a window size W ⁴. Each arm follows a
200 distribution with mean $\mu_i \in [0, 1]$. The arms arrive one by one in the stream, and we let $\{\text{arm}_i\}_{i=t-W+1}^t$ be
201 the set of valid arms that arrived in the W latest steps. When a new arm arrives, the algorithm can pull the
202 arriving arm and the arms in memory. The algorithm can also decide whether to store the new arm in memory
203 or discard it, and the algorithm can discard some arms stored in memory to free up space. At any point, the
204 collection of arms that the algorithm could access are the arms in memory and the arriving arm.

205 **Remark 1.** To keep consistent with the literature in sliding-window streaming algorithms, e.g., Datar et al.
206 (2002); Datar & Motwani (2016), we do *not* force the algorithm to discard the expired arms. Nevertheless,
207 our upper and lower bounds in Result 1 and Result 2 do *not* rely on this property. In other words, if we add
208 the condition that the expired arms have to be deleted immediately, the upper and lower bounds in Result 1
209 and Result 2 still hold. The immediate removal of expired arms from the memory is helpful for applications
210 with private data retention requirements.

211 We can now define the *sample and space complexity* of a sliding-window streaming MABs algorithm.

212 **Definition 3** (*Sample complexity*). The *sample complexity* of a sliding-window streaming MABs algorithm is
213 defined as the total number of pulls of the algorithm.

214 **Definition 4** (*Space complexity*). The *space complexity* of a sliding-window streaming algorithm is defined
215 as the maximum number of arms that we store in the memory at any time during the algorithm.

216 **Pure exploration.** One of the most natural problems in the MABs problem in the sliding-window model
217 is the *pure exploration* problem, where the algorithm is asked to return the best or near-best arms. In what
218 follows, we discuss the necessary notions before formally defining the pure exploration problems.

219 **Definition 5** (*Best arm in the window*). Assume that we have a collection of n arms $\{\text{arm}_i\}_{i=1}^n$ with means
220 μ_i and arranged in the streaming arriving ordered. Let W be the window size and t be the index of the current
221 arriving arm. Then, for any $t \in [n]$, the best arm in the window $\text{arm}^*(W, t)$ is the arm with the highest mean
222 $\mu^*(W, t)$ among the W latest arms $\{\text{arm}_i\}_{i=t-W+1}^t$.

223 Note that the notation $\text{arm}^*(W, t)$ is a function of t and W . We also call the set of arms $\{\text{arm}_i\}_{i=t-W+1}^t$
224 *valid* at time step t for fixed t and W .

225 We are ready to introduce the *pure exploration* problem for the sliding-window streaming MABs model.

226 **Problem 1** (Exact pure exploration in sliding-window MABs). Given a stream of n arms $\{\text{arm}_i\}_{i=1}^n$ and a
227 window size W , we say a sliding-window streaming MABs algorithm ALG solves

228

229 - *weak* pure exploration with probability $1 - \delta$ if at any time $t \in [n]$, ALG can output the best arm in the
230 window with probability at least $1 - \delta$.

231

232 ⁴We emphasize that the parameter W is an input parameter (not the algorithm's choice).

235 • *strong* pure exploration with probability $1 - \delta$ if ALG can output the best arm in the window at all time
 236 $t \in [n]$ with probability $1 - \delta$.
 237

238 Next, we could analogously define the ε exploration problem in both the *weak* and the *strong* versions for the
 239 sliding-window streaming MABs.

240 **Problem 2** (ε exploration in sliding-window MABs). Given a stream of n arms $\{\text{arm}_i\}_{i=1}^n$, a window size
 241 W , and a parameter ε , we say a sliding-window streaming MABs algorithm ALG solves

242 • *weak* ε exploration with probability $1 - \delta$ if at any time $t \in [n]$, ALG is able to output an arm with mean
 243 reward μ such that $\mu \geq \mu^*(t, W) - \varepsilon$ with probability at least $1 - \delta$.
 244 • *strong* ε exploration with probability $1 - \delta$ if ALG is able to output an arm with mean reward μ such that
 245 $\mu \geq \mu^*(t, W) - \varepsilon$ at all time $t \in [n]$ with probability $1 - \delta$.

246 Here, as defined in Definition 5, $\mu^*(t, W)$ is the mean reward of the best arm in the window.
 247

248 **Regret minimization.** In Section 1.1, we have discussed the high-level definition for our regret notion in
 249 sliding windows, i.e., the epoch-wise regret. We now introduce the formal definition as follow.

250 **Definition 6** (Regret minimization with epoch-wise regrets). Let $\{\text{arm}_i\}_{i=1}^n$ be a collection of n arms, and
 251 let W and T be the window size and the total number of trials. We divide T into $(n - W + 1)$ equal-sized
 252 *epochs* with $\frac{T}{n-W+1}$ in each epoch. Let t be the variable for the index of the arriving arm, and for any t , the
 253 algorithm is required to conduct *exactly* $\frac{T}{n-W+1}$ arm pulls among $\{\text{arm}_i\}_{i=t-W+1}^t$. We define the regret
 254 of the j -th epoch as $R^E(j) = \sum_{\tau=1}^{T/(n-W+1)} (\text{arm}^*(W, t) - \text{arm}_{i(\tau)})$, where $i(\tau)$ is the arm index pulled by
 255 the algorithm. The total regret is defined as $R_T = \sum_{j=1}^{T/(n-W+1)} R^E(j)$, i.e., the regret over the epochs.
 256

257 3 A LOWER BOUND FOR PURE EXPLORATION IN SLIDING-WINDOW MABs

260 The most natural pure exploration problem is *pure exploration* which asks to return the *best arm*. In the vanilla
 261 streaming multi-armed bandits (MABs) model, pure exploration can be solved with $O(n/\Delta_{[2]}^2)$ samples and
 262 a single-arm memory, where $\Delta_{[2]}$ represents the difference between the mean of the best and the second-best
 263 arms. As such, one would naturally wonder whether the same story applies to the sliding-window model. In
 264 this section, we will show that pure exploration is surprisingly much harder in the sliding-window streams:
 265 unless the algorithm uses $\Omega(W)$ space, we cannot obtain any algorithm that solves pure exploration.

266 The hard instance for our lower bound is a stream with descending mean rewards of arms, i.e., $\mu_1 > \mu_2 >$
 267 $\dots > \mu_n$ for arms . The optimal solution for the sliding-window MABs would be to select arm_{n-W+1} ,
 268 which is the oldest non-expired arm. However, to always keep the oldest arm that has not expired in the
 269 memory, we would naturally need W memory. The following theorem formalizes the above intuitions.

270 **Theorem 1.** *Any algorithm that given n arms in a sliding-window stream with a window size of W , solves the
 271 weak or strong pure exploration problem in sliding-window streaming multi-armed bandits with a probability
 272 of at least 99/100 has a space complexity of at least $\Omega(W)$, even if the sample complexity is unbounded.*

273 *Proof.* We prove the theorem for weak pure exploration, since the task of strong pure exploration is only
 274 harder. In other words, since the answer for strong exploration is always valid for weak exploration, the
 275 former task should use at least the same amount of memory and samples.
 276

277 By Yao’s minimax principle (Yao, 1977), it is sufficient to prove the lower bound for deterministic algorithms
 278 over a challenging distribution of inputs. Let $n = 2W$. We construct the instance $\{\text{arm}_i\}_{i=1}^n$ such that
 279 $\mu_i = 1 - \frac{i}{3W}$. To solve the *weak* pure exploration problem with a probability of at least $\frac{99}{100}$, the algorithm
 280 must correctly identify at least $\frac{49}{50}$ of the best arms in the second half of the stream $\{\text{arm}_i\}_{i=1}^n$. If the algorithm
 281 fails to do this, the overall success probability would drop below $1 \cdot \frac{1}{2} + \frac{49}{50} \cdot \frac{1}{2} = \frac{99}{100}$.

Let $T \subset \{W+1, W+2, \dots, 2W\}$ represent the collection of times when the algorithm correctly identifies the best arm in the window during the second half of the stream. Define $A = \{\text{arm}^*(W, t) | t \in T\}$ as the set of best arms in the window at times $t \in T$. For any $t \in \{W+1, W+2, \dots, 2W\}$, the best arm in the window $\text{arm}^*(W, t)$ should be arm_{t-W+1} because the expected values of the arms monotonically decrease in this instance. Therefore, we have $A = \{\text{arm}_{t-W+1} | t \in T\}$. Given that $T \subset \{W+1, W+2, \dots, 2W\}$ and $|T| \geq \frac{49}{50}W$, it follows that $A \subset \{\text{arm}_2, \text{arm}_3, \dots, \text{arm}_{W+1}\}$ and $|A| = |T| \geq \frac{49}{50}W$.

For any $W+1 \leq t < 2W$, $\text{arm}^*(W, t) = \text{arm}_{t-W+1}$ has already arrived by time $W+1$. Therefore, for any $t \in T \cap [2W-1]$, $\text{arm}^*(W, t)$ must be stored in memory by time $W+1$ so that it can be returned at time t . This means that at least $|A| - 1 = \frac{49}{50}W - 1$ arms must be stored in memory at time $W+1$. Hence, according to Yao's minimax principle, the algorithm must have a *space complexity* of at least $\Omega(W)$. \square

Note that the success probability of 99/100 in the theorem is not inherently special: by a simple probability boosting argument, we can always maintain $O(1)$ copies of the algorithm and output the majority with asymptotically the same memory and number of samples. As such, our lower bound in Theorem 1 applies to any success probability of $1/2 + \Omega(1)$.

4 SLIDING-WINDOW ALGORITHMS AND LOWER BOUNDS FOR ε -PURE EXPLORATION

Section 3 depicts a very pessimistic picture for the pure exploration of the *best arm* in sliding-window streaming MABs. A natural question to follow is whether we could get positive results using a relaxed notion. A natural candidate for this purpose is the ε exploration under the (ε, δ) -PAC framework. Here, instead of returning the *single best* arm, we are allowed to obtain an arm whose gap is within ε additive to the best, i.e., return an arm with mean reward $\mu \geq \mu^* - \varepsilon$. In this section, we present the bounds for both *strong* and *weak* ε exploration. Our main results are:

- A pure exploration algorithm that solves *weak* ε exploration with probability $1 - \delta$ in the sliding-window streaming MABs model with $O\left(\frac{n}{\varepsilon^2} \log \frac{W}{\delta}\right)$ sample complexity and $O\left(\frac{1}{\varepsilon}\right)$ space complexity.
- A lower bound shows that for any algorithm to solve *strong* ε exploration with probability 99/100 in the sliding-window streaming MABs model, the algorithm has to use $\Omega\left(\frac{n}{\varepsilon^2} \log \frac{n}{W}\right)$ sample complexity. Since $n \gg W$ in most cases, the lower bound separated the *weak* and *strong* ε exploration problems in the sliding-window streaming MABs model.
- Finally, we give a nearly-matching algorithm for *strong* ε exploration with probability $1 - \delta$ in the sliding-window streaming MABs model with $O\left(\frac{n}{\varepsilon^2} \log \frac{n}{\delta}\right)$ sample complexity and $O\left(\frac{1}{\varepsilon}\right)$ space complexity.

4.1 AN EFFICIENT ALGORITHM FOR WEAK ε -PURE EXPLORATION

We start with introducing a streaming algorithm designed for *weak* ε exploration.

Theorem 2. *There exists a streaming algorithm that, given n arms arriving in a sliding-window stream with a window size W and a confidence parameter δ , solves weak ε exploration with a probability of at least $1 - \delta$ using a sample complexity of $O\left(\frac{n}{\varepsilon^2} \log \frac{W}{\delta}\right)$ and a space complexity of $O\left(\frac{1}{\varepsilon}\right)$.*

At a high level, the algorithm follows the idea of partitioning the range $[0, 1]$ into $O\left(\frac{1}{\varepsilon}\right)$ segments (“buckets”) of equal length. An arm is considered to belong to a bucket if its mean value falls within the range of that segment. For an arm arm_i that belongs to bucket B , any arm arm' that is in a nearby bucket would serve as an ε -approximation of arm_i . If we pull each arm an adequate number of times, we can ensure that any arm is placed into a bucket that is close enough to its mean; thus, the non-expired arm from the highest bucket will be an ε -best arm. To optimize memory usage, we store only the latest arm for each bucket instead of all the arms that belong to that bucket. Our algorithm for *weak* ε exploration is presented in Algorithm 1, with the pulling size set to $s = (9/2\varepsilon^2) \cdot \ln 6W/\delta$.

329 **Algorithm 1:** Efficient Algorithm for ε exploration in Sliding-window Streaming MABs: $\text{BUCKET}(s)$

330 **Input:** Data stream $\{\text{arm}_i\}_{i=1}^n$, window size W , confidence parameter δ and accuracy parameter ε ;

331 **Input:** Sample complexity: $s = \frac{9}{2\varepsilon^2} \ln \frac{6W}{\delta}$ for weak exploration and $s = \frac{9}{2\varepsilon^2} \ln \frac{6n}{\delta}$ for strong

332 exploration;

333 **Output:** ε -best arms $\{\widehat{\text{arm}}_i\}_{i=1}^n$;

334 $N \leftarrow \frac{3}{\varepsilon}$;

335 Generate N buckets B_1, B_2, \dots, B_N ;

336 **for** each arriving arm arm_i **do**

337 Pull arm_i for s times and evaluate empirical mean $\widehat{\mu}_i$;

338 Store arm_i in B_j such that $(j-1)\frac{\varepsilon}{3} < \widehat{\mu}_i \leq j\frac{\varepsilon}{3}$ and discard the arms stored in B_j previously;

339 Discard all stored arms that are expired;

340 $\widehat{\text{arm}}_i \leftarrow$ the arm stored in B_k such that $k = \max_{i \leq N} \{B_i \neq \emptyset\}$

341 **end**

342 **return** $\{\widehat{\text{arm}}_i\}_{i=1}^n$

344

345

346 4.2 A LOWER BOUND FOR STRONG ε -PURE EXPLORATION

347

348 We will now discuss the lower bound for *strong ε exploration* that has an extra $\log n$ factor. In particular,

349 if we show that when $W \ll n$ (e.g., $W = \log n$), there is a lower bound of $\Omega(\frac{n}{\varepsilon^2} \log n)$ samples for strong

350 exploration, it would imply a *separation* between the *weak* and *strong ε exploration* since the weak exploration

351 only requires $\Omega(\frac{n}{\varepsilon^2} \log W)$ samples by Algorithm 1 $\text{BUCKET}(\frac{9}{2\varepsilon^2} \ln \frac{6W}{\delta})$. Then we have:

352 **Lemma 4.1.** *For infinitely many choices of parameters n, ε , and $W \leq n^{0.99}$, there exists a distribution of*

353 *arms $\mathcal{D}(n, W, \varepsilon)$ such that any algorithm that solves the strong ε exploration with probability at least 99/100*

354 *on $\mathcal{D}(n, W, \varepsilon)$ requires at least $\Omega(\frac{n}{\varepsilon^2} \log n)$ samples. The lower bound holds even if the algorithm is with*

355 *unbounded memory.*

356 The technical statement for Lemma 4.1 is more general and gives $\Omega(\frac{n}{\varepsilon^2} \log \frac{n}{W})$ samples for $W \in [1, n/8]$,

357 although the bound is less informative when W is large. At a high level, our lower bound works by reducing

358 solving *independent* copies of the ε -best arm identification to the sliding-window streaming ε exploration

359 case. Mannor & Tsitsiklis (2004) proved that $O(\frac{n}{\varepsilon^2} \log (\frac{1}{\delta}))$ pulls are necessary to identify an ε -best arm

360 among n arms with a probability of at least $1 - \delta$.

361 In the slide-window setting, since arms will expire after W time, the information from one window does

362 not affect another *disjoint* window. There are $\Theta(\frac{n}{W})$ windows in a sliding-window stream that are disjoint.

363 Since each window requires at least $O(\frac{W}{\varepsilon^2} \log (\frac{n}{W}))$ pulls to solve its exploitation with a probability of at

364 least $1 - \Theta(\frac{W}{n})$, it follows that $O(\frac{n}{\varepsilon^2} \log (\frac{n}{W}))$ pulls are necessary to achieve *strong ε exploration* with a

365 probability of at least 99/100.

366

367

368 4.3 AN EFFICIENT ALGORITHM FOR STRONG ε -PURE EXPLORATION

369 We introduce a streaming algorithm for *strong ε exploration*. The algorithm uses essentially the same

370 subroutine as in Algorithm 1, but it uses a larger pulling size of $s = \frac{9}{2\varepsilon^2} \ln \frac{6n}{\delta}$ to beat a union bound.

371 **Lemma 4.2.** *There exists a streaming algorithm that, given n arms arriving in a sliding-window stream with*

372 *a window size W and a confidence parameter δ , solves strong ε exploration with a probability of at least*

373 *$1 - \delta$. This algorithm achieves a sample complexity of $O(\frac{n}{\varepsilon^2} \log \frac{n}{\delta})$ and a space complexity of $O(\frac{1}{\varepsilon})$.*

376

5 REGRET MINIMIZATION IN SLIDING-WINDOW STREAMING MABS

377

378 In this section, we investigate *regret minimization* for sliding-window streaming multi-armed bandits (MABs).
379 Recall that in Definition 6, we defined regret minimization with the concepts of *epoch-wise* regret. Here, we
380 have $n - W + 1$ equal-sized epochs, and we must perform $\frac{T}{n-W+1}$ pulls in each epoch. The question is how
381 to minimize the cumulative regret over the entire horizon $[T]$.
382

383 The most natural idea is to adapt strategies in streaming MABs, e.g., (Wang, 2023), to get a low regret
384 algorithm. In particular, when a new arm arrives, we can use Algorithm 1 to pull the arm $O(\frac{1}{\varepsilon^2} \log n)$ times
385 and place it in the bucket. By the guarantees of Algorithm 1, we will be able to get ε -best arms at any step
386 with high probability. This strategy incurs a regret of $O(\frac{1}{\varepsilon^2} \log n)$ when identifying the ε -best arm during
387 each epoch. Additionally, there is a regret of $O(\varepsilon \frac{T}{n-W+1})$ for the remaining pulls on the ε -best arm we
388 identify within each epoch. As a result, the total regret is $O(\frac{n}{\varepsilon^2} \log n + \varepsilon T)$. The regret is minimized by
389 choosing $\varepsilon = O(\sqrt[3]{\frac{n \log n}{T}})$, which gives a total regret of $O(T^{\frac{2}{3}}(n \log n)^{\frac{1}{3}})$.
390

391 Alas, this strategy has a fatal issue: Algorithm 1 requires $O(\frac{1}{\varepsilon})$ memory space; and since in most cases
392 $T \gg n \gg W$, the memory of $1/\varepsilon = O(\sqrt[3]{\frac{T}{n \log n}})$ could be way bigger than the window size W . Thus, it is
393 not immediately clear whether we could get low-regret algorithms with small memory in this setting. In this
394 section, we show that the issue of the aforementioned algorithm is not an artifact: we prove a strong lower
395 bound showing that a total regret of $O(\frac{T}{W^2})$ is unavoidable if we only have $o(W)$ space.
396

397 **Theorem 3.** *There exists a family of streaming stochastic multi-armed bandit instances such that, for any
398 given parameters T , n , and W , where $T \geq n \geq 16W$, any single-pass streaming algorithm for a sliding-
399 window stream of length n with a window size W and a memory capacity of $\frac{W-1}{2}$ arms must incur a total
400 expected regret given by $\mathbb{E}[R_T] \geq \frac{T}{64W^2}$.*
401

402 *Furthermore, there exists an algorithm that given n arms arriving in a stream and parameters W and T ,
403 achieves $O(\sqrt{W \cdot (n - W) \cdot T})$ total regret with W memory.*
404

405 At a high level, our lower bound is obtained by constructing W arms whose means decrease by $\frac{1}{W}$ and W
406 arms with the same mean, and the pattern is repeated over the stream. Since we can only store at most half of
407 these arms, if the best arm in the epoch is missed, the regret for each pull will be at least $\frac{1}{2W}$. This leads to a
408 total regret of $\Omega(\frac{T}{W^2})$. Our upper bound is obtained by running UCB-based algorithms on each window.
409

410

6 EXPERIMENTS

411

412 We conduct experiments for both ε -exploration and regret minimization in the sliding-window streaming
413 setting. Our main empirical finding is that, consistent with our theoretical results, both the ε -exploration
414 and regret minimization algorithms demonstrate trade-offs between memory and quality/regret. The regret
415 minimization algorithm demonstrates a sharp change around the $O(W)$ -arm memory. We will briefly
416 demonstrate the experiments of the ε -exploration and regret minimization algorithms in epoch-wise settings.
417 Additional experimental results can be found in ??.
418

419 **The data.** We use synthetic data with streams of arms to conduct our experiments. We use different types of
420 instances for exploration and regret minimization as follows.
421

422

- For exploration, we sample n arms with the distribution $\text{Bern}(p)$ such that p is from a uniform distribution⁵.
423 We note that the “uniform” type of instances are more suitable for ε -exploration since the quality decrement

424

425 ⁵We use $\text{Bern}(p)$ to denote Bernoulli distribution with mean p .
426

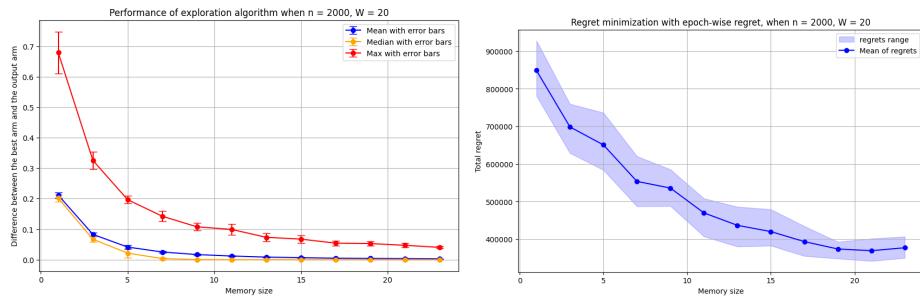
423 of the returned arms could be better captured. We use $n \in \{1000, 2000, 5000\}$ and $W \in \{10, 20, 50\}$ for
 424 ε -exploration experiments.

425 • For regret minimization, we need instance distributions *consistent* with our instance distribution in Section 5.
 426 For the epoch-wise regret minimization, we sample $n - n/W$ arms with distribution $\text{Bern}(0.25)$ and n/W
 427 arms with distribution $\text{Bern}(0.95)$. We then permute the arms uniformly. Due to constraints on running
 428 time, we use $n \in \{500, 1000, 2000\}$ and $W \in \{10, 20, 50\}$ for ε -exploration experiments.

429 To mitigate the noise from randomness, for each parameter setting with fixed memory size, we conduct
 430 10 **independent runs of experiments and take the average**. For the quality of the arm and the regret
 431 minimization, we also report error bars and the ranges of the regrets.

432 **The algorithms.** We conduct experiments for both exploration and regret minimization. We first implement
 433 and test our ε -exploration algorithm (Algorithm 1). **Note that it is hard to compare performances with baseline**
 434 **algorithms for sliding-window exploration since we cannot easily quantify “error” when the algorithm outputs**
 435 **an expired arm. Therefore, we report the performance of our algorithm with different memory sizes.**

436 For regret minimization, we adapt the algorithm with W -arm memory discussed in Section 5. To handle the
 437 case of $m < W$ -arm memory, we simulate the reservoir sampling: after the memory is full, for each arriving
 438 arm, we toss a fair coin with bias m/t for the t -th arriving arm to decide whether we admit the new arm to
 439 the memory (by uniformly at random discarding an arm existing in the memory).



451
 452 Figure 1: The performances of ε -exploration and regret minimization, $n = 2000$, $W = 20$.

453 **Summary of the experiments.** A sample of the performances for ε -exploration and regret minimization
 454 is given in Figure 1 (for $n = 2000$ and $W = 20$; see ?? for more parameter settings). As we can observe
 455 from the figures, for all the experiments, there is generally a trade-off between the arm quality/regret and
 456 memory. The trade-off in ε -exploration is generally smoother, and the regret minimization for the everlasting
 457 best arm demonstrates a sharp drop of regret around the W -memory point. These results are consistent with
 458 our theoretical findings for sliding-window streaming MABs algorithms.

461 7 CONCLUSION AND FUTURE WORK

462 In this work, we initiated the study of multi-armed bandits (MABs) in the sliding-window model. Our results
 463 built the fundamental hardness of online learning in the sliding-window MABs model, and we provided
 464 important insights for related applications, e.g., using ε -exploration rather than pure exploration in practice.
 465 There are several open directions to follow up on our work. For instance, one appealing question is the
 466 *multi-pass* setting: if the algorithm is allowed to make multiple passes over the stream, it might be possible
 467 for the algorithm to achieve better memory efficiency. The sliding-window model for other variants of MABs,
 468 e.g., the linear bandits, can be another interesting direction to pursue.

470 REFERENCES
471

472 Arpit Agarwal, Sanjeev Khanna, and Prathamesh Patil. A sharp memory-regret trade-off for multi-pass
473 streaming bandits. In Po-Ling Loh and Maxim Raginsky (eds.), *Conference on Learning Theory, 2-5 July*
474 *2022, London, UK*, volume 178 of *Proceedings of Machine Learning Research*, pp. 1423–1462. PMLR,
475 2022. URL <https://proceedings.mlr.press/v178/agarwal22a.html>.

476 Apple Inc. Differential privacy overview. Apple, 2021. URL https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf.

477 Sepehr Assadi and Chen Wang. Exploration with limited memory: streaming algorithms for coin tossing, noisy
478 comparisons, and multi-armed bandits. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani,
479 Gautam Kamath, and Julia Chuzhoy (eds.), *Proceedings of the 52nd Annual ACM SIGACT Symposium on*
480 *Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pp. 1237–1250. ACM, 2020. doi:
481 10.1145/3357713.3384341. URL <https://doi.org/10.1145/3357713.3384341>.

482 Sepehr Assadi and Chen Wang. Single-pass streaming lower bounds for multi-armed bandits exploration
483 with instance-sensitive sample complexity. In *NeurIPS*, 2022.

484 Sepehr Assadi and Chen Wang. The best arm evades: Near-optimal multi-pass streaming lower bounds for pure
485 exploration in multi-armed bandits. In Shipra Agrawal and Aaron Roth (eds.), *The Thirty Seventh Annual*
486 *Conference on Learning Theory, June 30 - July 3, 2023, Edmonton, Canada*, volume 247 of *Proceedings*
487 *of Machine Learning Research*, pp. 311–358. PMLR, 2024. URL <https://proceedings.mlr.press/v247/assadi24a.html>.

488 Michele Borassi, Alessandro Epasto, Silvio Lattanzi, Sergei Vassilvitskii, and Morteza Zadimoghaddam.
489 Sliding window algorithms for k-clustering problems. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia
490 Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing*
491 *Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December*
492 *6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/631e9c01c190fc1515b9fe3865abbb15-Abstract.html>.

493 Djallel Bouneffouf, Irina Rish, Guillermo A. Cecchi, and Raphaël Féraud. Context attentive bandits: Context-
494 tual bandit with restricted context. In Carles Sierra (ed.), *Proceedings of the Twenty-Sixth International*
495 *Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pp.
496 1468–1475. ijcai.org, 2017. doi: 10.24963/IJCAI.2017/203. URL <https://doi.org/10.24963/ijcai.2017/203>.

497 Vladimir Braverman and Rafail Ostrovsky. Smooth histograms for sliding windows. In *48th Annual IEEE*
498 *Symposium on Foundations of Computer Science (FOCS 2007), October 20-23, 2007, Providence, RI,*
499 *USA, Proceedings*, pp. 283–293. IEEE Computer Society, 2007. doi: 10.1109/FOCS.2007.55. URL
500 <https://doi.org/10.1109/FOCS.2007.55>.

501 Vladimir Braverman, Harry Lang, Keith D. Levin, and Morteza Monemizadeh. Clustering problems on
502 sliding windows. In Robert Krauthgamer (ed.), *Proceedings of the Twenty-Seventh Annual ACM-SIAM*
503 *Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pp. 1374–
504 1390. SIAM, 2016. doi: 10.1137/1.9781611974331.CH95. URL <https://doi.org/10.1137/1.9781611974331.ch95>.

505 Deepayan Chakrabarti, Ravi Kumar, Filip Radlinski, and Eli Upfal. Mortal multi-armed bandits. *Advances in*
506 *neural information processing systems*, 21, 2008.

517 Vincent Cohen-Addad, Shaofeng Jiang, Qiaoyuan Yang, Yubo Zhang, and Samson Zhou. Fair clustering
 518 in the sliding window model. In *Proceedings of the Thirteenth International Conference on Learning*
 519 *Representations (ICLR 2025, to appear)*, 2025.

520

521 Michael Crouch and Daniel S. Stubbs. Improved streaming algorithms for weighted matching, via
 522 unweighted matching. In *Approximation, Randomization, and Combinatorial Optimization. Algo-*
 523 *rithms and Techniques, APPROX/RANDOM 2014, September 4-6, 2014*, pp. 96–104, 2014. doi:
 524 10.4230/LIPIcs.APPROX-RANDOM.2014.96.

525 Michael S. Crouch, Andrew McGregor, and Daniel M. Stubbs. Dynamic graphs in the sliding-window
 526 model. In Hans L. Bodlaender and Giuseppe F. Italiano (eds.), *Algorithms - ESA 2013 - 21st Annual*
 527 *European Symposium, Sophia Antipolis, France, September 2-4, 2013. Proceedings*, volume 8125 of
 528 *Lecture Notes in Computer Science*, pp. 337–348. Springer, 2013. doi: 10.1007/978-3-642-40450-4_29.
 529 URL https://doi.org/10.1007/978-3-642-40450-4_29.

530

531 Mayur Datar and Rajeev Motwani. The sliding-window computation model and results. In Minos N.
 532 Garofalakis, Johannes Gehrke, and Rajeev Rastogi (eds.), *Data Stream Management - Processing High-*
 533 *Speed Data Streams*, Data-Centric Systems and Applications, pp. 149–165. Springer, 2016. doi: 10.1007/
 534 978-3-540-28608-0_7. URL https://doi.org/10.1007/978-3-540-28608-0_7.

535 Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Maintaining stream statistics over
 536 sliding windows. *SIAM J. Comput.*, 31(6):1794–1813, 2002. doi: 10.1137/S0097539701398363. URL
 537 <https://doi.org/10.1137/S0097539701398363>.

538 Alessandro Epasto, Mohammad Mahdian, Vahab S. Mirrokni, and Peilin Zhong. Improved sliding window
 539 algorithms for clustering and coverage via bucketing-based sketches. In Joseph (Seffi) Naor and Niv
 540 Buchbinder (eds.), *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022,*
 541 *Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pp. 3005–3042. SIAM, 2022. doi:
 542 10.1137/1.9781611977073.117. URL <https://doi.org/10.1137/1.9781611977073.117>.

543

544 Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the
 545 multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7:
 546 1079–1105, 2006.

547 GDPR. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the
 548 protection of natural persons with regard to the processing of personal data and on the free movement
 549 of such data, and repealing directive 95/46/ec (general data protection regulation), 2016. URL <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

550

551 Google LLC. How google retains data we collect. Google, 2025. URL <https://policies.google.com/technologies/retention?hl=en-US>. Accessed January 29, 2025.

552

553

554 Francesco Gullo, Domenico Mandaglio, and Andrea Tagarelli. A combinatorial multi-armed ban-
 555 dit approach to correlation clustering. *Data Min. Knowl. Discov.*, 37(4):1630–1691, 2023. doi:
 556 10.1007/S10618-023-00937-5. URL <https://doi.org/10.1007/s10618-023-00937-5>.

557

558 Yuchen He, Zichun Ye, and Chihao Zhang. Understanding memory-regret trade-off for streaming stochastic
 559 multi-armed bandits. In *Proceedings of the 2025 ACM-SIAM Symposium on Discrete Algorithms, SODA*
 560 *2025*, 2025. doi: 10.48550/ARXIV.2405.19752. URL <https://doi.org/10.48550/arXiv.2405.19752>.

561

562 Tianyuan Jin, Keke Huang, Jing Tang, and Xiaokui Xiao. Optimal streaming algorithms for multi-armed
 563 bandits. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference*

564 *on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event, volume 139 of Proceedings of Machine Learning Research, pp. 5045–5054. PMLR, 2021. URL <http://proceedings.mlr.press/v139/jin21a.html>.*

565

566

567 Zohar Shay Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013, volume 28 of JMLR Workshop and Conference Proceedings, pp. 1238–1246. JMLR.org, 2013. URL <http://proceedings.mlr.press/v28/karnin13.html>.*

571

572 Nikolai Karpov and Chen Wang. Nearly tight bounds for exploration in streaming multi-armed bandits with known optimality gap. In Toby Walsh, Julie Shah, and Zico Kolter (eds.), *AAAI-25, Sponsored by the Association for the Advancement of Artificial Intelligence, February 25 - March 4, 2025, Philadelphia, PA, USA, pp. 17788–17796. AAAI Press, 2025. doi: 10.1609/AAAI.V39I17.33956. URL <https://doi.org/10.1609/aaai.v39i17.33956>.*

573

574

575

576

577 David Liau, Zhao Song, Eric Price, and Ger Yang. Stochastic multi-armed bandits in constant space. In Amos J. Storkey and Fernando Pérez-Cruz (eds.), *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain, volume 84 of Proceedings of Machine Learning Research, pp. 386–394. PMLR, 2018. URL <http://proceedings.mlr.press/v84/liau18a.html>.*

578

579

580

581

582 David E Losada, Javier Parapar, and Alvaro Barreiro. Multi-armed bandits for adjudicating documents in pooling-based evaluation of information retrieval systems. *Information Processing & Management*, 53(5):1005–1025, 2017.

583

584

585

586 Arnab Maiti, Vishakha Patil, and Arindam Khan. Multi-armed bandits with bounded arm-memory: Near-optimal guarantees for best-arm identification and regret minimization. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pp. 19553–19565, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/a2f04745390fd6897d09772b2cd1f581-Abstract.html>.*

587

588

589

590

591

592

593 Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.

594

595

596 Philip Pallmann, Alun W Bedding, Babak Choodari-Oskooei, Munyaradzi Dimairo, Laura Flight, Lisa V Hampson, Jane Holmes, Adrian P Mander, Lang’o Odondi, Matthew R Sydes, et al. Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC medicine*, 16:1–15, 2018.

597

598

599 Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In William W. Cohen, Andrew McCallum, and Sam T. Roweis (eds.), *Machine Learning, Proceedings of the Twenty-Fifth International Conference (ICML 2008), Helsinki, Finland, June 5-9, 2008, volume 307 of ACM International Conference Proceeding Series, pp. 784–791. ACM, 2008. doi: 10.1145/1390156.1390255. URL <https://doi.org/10.1145/1390156.1390255>.*

600

601

602

603

604 Herbert Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 55:527–535, 1952.

605

606

607 Weiwei Shen, Jun Wang, Yu-Gang Jiang, and Hongyuan Zha. Portfolio choices with orthogonal bandit learning. In Qiang Yang and Michael J. Wooldridge (eds.), *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015, pp. 974. AAAI Press, 2015. URL <http://ijcai.org/Abstract/15/142>.*

608

609

610

611 David Simchi-Levi and Chonghuan Wang. Multi-armed bandit experimental design: Online decision-making
 612 and adaptive inference. In Francisco J. R. Ruiz, Jennifer G. Dy, and Jan-Willem van de Meent (eds.),
 613 *International Conference on Artificial Intelligence and Statistics, 25-27 April 2023, Palau de Congressos,*
 614 *Valencia, Spain*, volume 206 of *Proceedings of Machine Learning Research*, pp. 3086–3097. PMLR, 2023.
 615 URL <https://proceedings.mlr.press/v206/simchi-levi23a.html>.

616 Yufei Tao and Dimitris Papadias. Maintaining sliding window skylines on data streams. *IEEE Trans. Knowl.*
 617 *Data Eng.*, 18(2):377–391, 2006. doi: 10.1109/TKDE.2006.48. URL <https://doi.org/10.1109/TKDE.2006.48>.

618 Francesco Trovò, Stefano Paladino, Marcello Restelli, and Nicola Gatti. Improving multi-armed bandit
 619 algorithms in online pricing settings. *International Journal of Approximate Reasoning*, 98:196–235, 2018.

620 Chen Wang. Tight regret bounds for single-pass streaming multi-armed bandits. In Andreas Krause, Emma
 621 Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International*
 622 *Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202
 623 of *Proceedings of Machine Learning Research*, pp. 35525–35547. PMLR, 2023. URL <https://proceedings.mlr.press/v202/wang23a.html>.

624 Peter Whittle. Restless bandits: Activity allocation in a changing world. *Journal of applied probability*, 25
 625 (A):287–298, 1988.

626 David P. Woodruff, Peilin Zhong, and Samson Zhou. Near-optimal k-clustering in the sliding win-
 627 dows model. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and
 628 Sergey Levine (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference*
 629 *on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December*
 630 *10 - 16, 2023*, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/476ab8f369e489c04187ba84f68cfa68-Abstract-Conference.html.

631 Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity. *18th Annual*
 632 *Symposium on Foundations of Computer Science (sfcs 1977)*, pp. 222–227, 1977. URL <https://api.semanticscholar.org/CorpusID:143169>.

633 Chao Zhang, Angela Bonifati, and M. Tamer Özsu. Incremental sliding window connectivity over streaming
 634 graphs. *Proc. VLDB Endow.*, 17(10):2473–2486, 2024. doi: 10.14778/3675034.3675040. URL <https://www.vldb.org/pvldb/vol17/p2473-zhang.pdf>.

635 Liangwei Zhang, Jing Lin, and Ramin Karim. Sliding window-based fault detection from high-dimensional
 636 data streams. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(2):289–303, 2016.

637
 638
 639
 640
 641
 642
 643
 644
 645
 646
 647
 648
 649
 650
 651
 652
 653
 654
 655
 656
 657