Online Learning for Non-Stationary A/B Tests

Andrés Muñoz Medina
Google AI
New York, NY
ammedina@google.com

Sergei Vassilvitskii
Google AI
New York, NY
sergeiv@google.com

Dong Yin
Dept. of EECS, UC Berkeley
Berkeley, CA
dongyin@berkeley.edu

ABSTRACT

The rollout of new versions of a feature in modern applications is a manual multi-stage process, as the feature is released to ever larger groups of users, while its performance is carefully monitored. This kind of A/B testing is ubiquitous, but suboptimal, as the monitoring requires heavy human intervention, is not guaranteed to capture consistent, but short-term fluctuations in performance, and is inefficient, as better versions take a long time to reach the full population.

In this work we formulate this question as that of expert learning, and give a new algorithm *Follow-The-Best-Interval*, FTBI, that works in dynamic, non-stationary environments. Our approach is practical, simple, and efficient, and has rigorous guarantees on its performance. Finally, we perform a thorough evaluation on synthetic and real world datasets and show that our approach outperforms current state-of-the-art methods.

CCS CONCEPTS

• **Information systems** → **Computational advertising**; • **Theory of computation** → **Online learning algorithms**; **Computational pricing and auctions**; • **Security and privacy** → **Database activity monitoring**;

KEYWORDS

A/B tests; expert learning; non-stationary environments

ACM Reference Format:

Andrés Muñoz Medina, Sergei Vassilvitskii, and Dong Yin. 2018. Online Learning for Non-Stationary A/B Tests. In *The 27th ACM International Conference on Information and Knowledge Management (CIKM '18)*, October 22–26, 2018, Torino, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3269206.3271718>

1 INTRODUCTION

Whether it is a minor tweak, or a major new update, releasing a new version of a running system is a stressful time. While the release has typically gone through rounds of offline testing, real world testing often uncovers additional corner cases that may manifest themselves as bugs, inefficiencies, or overall poor performance. This is especially the case in machine learning applications, where models are typically trained to maximize a proxy objective, and a

model that performs better on offline metrics is not guaranteed to work well in practice.

The usual approach in such scenarios is to evaluate the new system through a series of closely monitored A/B tests. The new version is usually released to a small number of customers, and, if no concerns are found and metrics look good, the portion of traffic served by the new system is slowly increased.

While A/B tests provide a sense of safety in that a detrimental change will be quickly observed and corrected (or rolled back), they are not a silver bullet. First, A/B tests are labor intensive—they are typically monitored manually, with an engineer, or a technician, checking the results of the test on a regular basis (for example, daily or weekly). Second, the evaluation is usually dependent on average metrics—e.g. increasing clickthrough rate, decreasing latency, etc. However, fluctuations in that metric can be easily missed—for instance a system that performs well during the day, but lags at night, may appear better on a day-to-day view, and thus its sub-optimality at night is missed. Finally, due to their staged rollouts, A/B tests are inefficient—the better version of the system remains active only on a small subset of traffic until a human intervenes, even if it is universally better right off the bat.

We propose a method for automated version selection that addresses all of the inefficiencies described above. At a high level, we cede the decision of whether to use the old version or a new version to an automated agent. The agent repeatedly evaluates the performance of both versions and selects the better one automatically. As we describe below, this, by itself, is not a new approach, and is captured in the context of *expert* or *bandit learning* in the machine learning literature. However, the vast majority of previous work optimizes for the average case, where the distribution of real world examples is fixed (but unknown). This modeling choice cannot capture short term but consistent fluctuations in efficacy of one system over another. The main contribution of this work is a simple, practical, and efficient algorithm that selects the best version even under non-stationary assumptions; and comes with strong theoretical guarantees on its performance.

The paper is organized as follows: first we describe the problem of learning with expert advice and discuss previous work related to the contents of this paper. In Section 2 we formalize the expert learning scenario we will be analyzing. Our proposed algorithm and its performance analysis are presented in Sections 3 and 4 respectively. Finally, in Section 5 we conduct extensive empirical evaluations demonstrating the advantages of our algorithm over the current state of the art.

1.1 Expert Learning

Learning with expert advice is a classical setting in online learning. Over a course of T time steps, a player is repeatedly asked to decide on a best action in a particular setting. For example, in parallel

This work is licensed under a Creative Commons Attribution International 4.0 License.

CIKM '18, October 22–26, 2018, Torino, Italy
© 2018 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6014-2/18/10.
<https://doi.org/10.1145/3269206.3271718>

computing, the agent may need to decide how many workers to allocate for a particular computation; in advertising auctions, the auctioneer needs to decide on a minimum (reserve) price to charge; in ranking systems, the algorithm needs to decide which answer to rank at the top, and so on.

The player does not make these decisions in a vacuum, instead he has access to a number of *experts*. Each of the experts may provide her own answer to the query; in our setting each expert represents a different version of the system. At a high level, the goal of the player is to perform (almost) as well as the best expert overall. The difference between the performance of the player and the best expert is known as *regret*, thus the player's goal is equivalent to minimizing his regret. Note that if all experts perform poorly, the player is not penalized for not doing well, whereas if one expert does well, the player is best served by honing in on this 'true' expert and following her recommendations.

To prove bounds on the worst case regret achieved by the player following a particular strategy, we must make assumptions on the distribution from which queries are drawn. In the most studied versions of the problem, the environment is either stochastic (queries are drawn from a fixed, but unknown, distribution), or adversarial (where no assumptions are made). While it is not realistic to assume a fully stochastic setting, neither is an adversarial setting a good model for distributions observed in practice. Rather, the environment is better represented as a series of stochastic intervals, that may have sharp and unpredictable transitions between them. For a concrete example, consider the problem of predicting time to travel from point A to point B. If we have two experts, one performing well when the roads are congested, and another when the roads are relatively free, then a day may have multiple periods that are locally stochastic, but are drastically different from each other. Moreover, while some congested times are predictable, such as the morning and evening rush hour, those caused by accidents, special events, and holidays are not, and represent the scenario we target.

To model this setting, several notions such as *shifting regret* [12] and *adaptive regret* [11] have been introduced, and these works try to ensure that the performance of the algorithm over sub-intervals of time is comparable to the best expert for that time interval. Several algorithms were developed for this setting, and we provide a more comprehensive review of these algorithms in Section 1.2. Among these algorithms, AdaNormalHedge developed by Luo and Shapire [19] represents the current state of the art, since it is relatively fast, parameter-free, and gives provable regret bounds.

While AdaNormalHedge seems like the answer we seek, it has never been empirically validated and has some disadvantages which render it less effective in practice. The straightforward implementation has update time linear in the time elapsed since the start, and, while compression tricks borrowed from streaming literature reduce the bound to polylogarithmic in theory, these updates are nevertheless slow and costly in practice. However, the main drawback of this algorithm is that it is not intuitive, and thus is hard to debug and interpret.

In this work, we propose a different approach. Our algorithm, *Follow-The-Best-Interval* (FTBI) builds on the standard Follow-The-Leader (FTL) algorithm for purely stochastic environments, and is easy to explain. We prove (see Theorem 1) that in the case of two experts, the algorithm has regret polylogarithmic in T , and

quadratic in the number of changes in the environment. We also show that the update time complexity of our algorithm is $O(\log T)$, which is provably better than that of AdaptiveNormalHedge even with the streaming technique. The regret bound of our algorithm is slightly weaker than that given by AdaNormalHedge; however as we show through extensive simulation and real-data experimental results, in practice FTBI performs significantly better.

1.2 Related Work

Our setup best matches that of learning with expert advice. This area of machine learning has been well studied since the pioneering works of [4, 9, 18, 25]. There are two common assumptions on the nature of the rewards observed by the agent: stochastic and adversarial rewards. In the stochastic setting, Follow-The-Leader (FTL) is an easy and intuitive algorithm that can achieve constant regret; while for the adversarial setting, the Weighted Majority algorithm [9, 18] and the Follow the Perturbed Leader algorithm [14, 16] are the most commonly used algorithms with regret in $\tilde{O}(\sqrt{T})$.

The aforementioned algorithms measure regret against a single fixed best expert. This is a relatively simple and well-understood scenario. By contrast the non-stationary environments generally pose greater challenges in the design of regret minimizing algorithms. In this setting, the full learning period $[T]$ is usually partitioned into a few segments; and the algorithm aims to compete against the sequence of best experts over each segment. The regret in this setting is also known as *shifting regret* (See Section 2 for a precise definition). Similar to the non-adaptive scenario, we model the input as coming from stochastic or adversarial settings. In the stochastic setting, the rewards of each expert remain i.i.d. in each segment, but the distribution may change across the segments. For the adversarial setting, no assumptions are made as to how the rewards are generated. The first analysis of online learning with shifting environments was given by Herbster and Warmuth [12]. The authors proposed the Fixed-Share algorithm (and a few variants) to achieve low shifting regret in the adversarial setting. The idea of Fixed-Share, and other derivations of this algorithm [1–3, 6, 13], is to combine the Weighted Majority algorithm with a mixing strategy. In general these algorithms are guaranteed to achieve regret in $\tilde{O}(\sqrt{T})$ with an additional penalty for shifting.

A different framework for solving this problem is the so called sleeping expert technique, originally introduced by Freund et al. [10]. Although the algorithm achieves good shifting regret bounds, in its original form, this algorithm has prohibitively high computational costs. To address this issue, Hazan and Seshadhri [11] propose the Follow the Leading History (FLH) algorithm with a data streaming technique to reduce the time complexity of sleeping expert algorithms. The authors provide an algorithm with logarithmic shifting regret. However, this bound holds only for rewards distributed according to an exp-concave distribution. Since we make no assumption on the distribution generating our rewards, we cannot extend the results of FLH to our setup.

One final line of work dealing with shifting regret under adversarial rewards is the so called strongly adaptive online learning (SAOL) [8, 15, 27]. SAOL framework aims to achieve low regret over any interval, and thus is strictly harder than achieving low shifting regret. Despite this difficulty, there are algorithms that achieve regret in $O(\sqrt{T})$, matching the traditional shifting regret scenario.

Moreover, without any assumptions on the mechanism generating the rewards, this regret is in fact tight. In practice however, rewards are hardly fully adversarial. While we certainly don't expect rewards to be i.i.d. over the full time period, it is generally reasonable to assume we can model rewards as a stochastic process with shifts. For instance, rewards can be i.i.d. throughout a day but change on weekends or holidays. When this is the case, one can obtain exponential improvements and achieve regret of $O(\text{poly}(\log T))$. The design of an algorithm with this regret bound in the stochastic setting with shifts was proposed as an open problem by Warmuth et al. [26]. A few solutions have been given since then. Sani et al. [22] propose an algorithm that can achieve low shifting regret in both stochastic and adversarial settings, however, their algorithm is not parameter-free and it requires proper tuning and knowledge of the number of shifts. Luo and Schapire [19] introduced the AdaNormal-Hedge (ANH) algorithm, a parameter free approach that achieves shifting regret of $O(\frac{N}{\Delta} \log T)$, where N is the number of shifts and Δ captures the gap between the best and second best choices.

While this bound is strictly better than the regret bound we provide for FTBI, we will extensively demonstrate that in practice FTBI consistently outperforms ANH both in speed and accuracy. Moreover, FTBI is a much more intuitive and explainable algorithm since it is a simple generalization of the classic Follow The Leader approach (FTL). By contrast, the state of ANH depends on an streaming technique [11, 19] which offers no transparency in the decisions made by the algorithm.

Finally, there has been a recent attempt to use multi-armed bandit methods to replace A/B tests [23, 24]. This techniques however, assume that one only has access to data from one of the experiments which differ from our full information scenario.

2 SETUP

Let $T > 0$ denote a time horizon and $[T] = \{1, \dots, T\}$. We consider the problem of designing an automated agent that, at every time $t \in [T]$ chooses between two version of a system $x_t \in \{V_1, V_2\}$. (The setting easily extends to K versions, and we will explore the performance of our algorithm for $K > 2$ in Section 5.) After the choice is made, a reward vector $\mathbf{r}_t = (r_t(V_1), r_t(V_2)) \in [0, 1]^2$ is revealed and the agent obtains a reward $r_t(x_t)$. We consider a full information, as opposed to a bandit setup since in practice we can run simultaneous experiments that can provide us with the reward information for each version. The goal of the monitoring system is to maximize the expected cumulative reward $\mathbb{E}[\sum_{t=1}^T r_t(x_t)]$. For instance, V_1 and V_2 could be two versions of a network allocation algorithm and $\mathbf{r}_t$ corresponds to the average number of queries per second the network can handle. We will abuse our notation and denote by $r_t(k) = r_t(V_k)$ the reward of using version V_k .

We now describe the reward generation mechanism. Let $0 = \tau_0 < \tau_1 < \dots < \tau_N = T$ be a partition of $[T]$, where N denotes the total number of shifts. For every $i \in [N]$, let $\mathcal{T}_i = [\tau_{i-1} + 1, \tau_i]$ be the i -th segment of the partition. We assume that for $t \in \mathcal{T}_i$, reward vectors $\mathbf{r}_t$ are drawn i.i.d according to an unknown distribution $\mathcal{D}_i$. We make no assumptions on $\mathcal{D}_i$ in particular $\mathcal{D}_i$ can differ from $\mathcal{D}_j$ for $i \neq j$.

Notice that if we consider each version to be an expert, we can cast our problem as that of learning with expert advice under *stochastic rewards* and *changing environments*. Therefore, borrowing

from the expert learning literature, we measure the performance of our agent using the shifting pseudo-regret.

Definition 1. For every $i \in [N]$, $t \in \mathcal{T}_i$ and $k \in \{1, 2\}$, let $\mu_i(k) = \mathbb{E}[r_t(k)]$ denote the expected reward of version V_k with respect to the distribution $\mathcal{D}_i$. Let $k_i^* = \arg\max_k \mu_i(k)$ and $\mu_i^* = \mu_i(k_i^*)$ denote the optimal version over segment $\mathcal{T}_i$ and its reward respectively. The shifting pseudo-regret is defined by

$$R_T := \sum_{i=1}^N \sum_{t=\tau_{i-1}+1}^{\tau_i} \mu_i^* - \mathbb{E}[r_t(x_t)],$$

where expectation is taken over both the reward vector and the randomization of the monitoring algorithm.

A successful algorithm for this scenario is one for which R_T is in $o(T)$. That is, the agent *learns* to choose the optimal version in a sublinear number of rounds.

Let $\Delta_i = \min_{k \neq k_i^*} \mu_i^* - \mu_i(k)$ denote the expected reward gap between the best and the next best version in $\mathcal{T}_i$, and let $\Delta = \min_{i \in [N]} \Delta_i$. Throughout the paper we will assume that $\Delta > 0$. Notice that this is without loss of generality as the case $\Delta = 0$ is uninteresting since playing either version would have yielded the same expected reward.

3 FOLLOW-THE-BEST-INTERVAL

We begin this section by describing the Follow-The-Leader (FTL) algorithm, a simple yet powerful algorithm for learning with expert advice under stochastic rewards (with no shifts). FTL maintains a weight $W_t(k) = \sum_{s=1}^t r_s(k)$ for each expert $k \in \{1, 2\}$. Weights correspond to the cumulative reward seen thus far. At time t , FTL chooses the expert with the highest weight, in other words, the leader. It is well known that if rewards are sampled i.i.d., FTL is guaranteed to achieve constant pseudo-regret in $O(\frac{1}{\Delta})$.

If we had access to the switching times $\tau_1, \dots, \tau_N$, we could simply run FTL and restart the weights kept by the algorithm at every breakpoint τ_j . This would guarantee pseudo-regret of $O(\frac{N}{\Delta})$. In the absence of this information, we could try to *detect* when a shift has occurred. However, it is not hard to see that, due to the noise in the rewards, detecting a shift accurately would require $O(\sqrt{T})$ time steps.

Instead, we propose the Follow-The-Best-Interval (FTBI) algorithm. At a high level, FTBI maintains a small number of FTL instances, each running over a different interval of a collection of nested subintervals of $[T]$. The algorithm then follows the action of the best performing FTL instance. More precisely, let

$$\mathcal{I}_n = \{[i \cdot 2^n, (i+1) \cdot 2^n - 1] : i \in \mathbb{N}\}.$$

Here, $\mathcal{I}_n$ is a set of disjoint intervals of length 2^n . Let $\mathcal{I} = \bigcup_{n=0}^{\infty} \mathcal{I}_n$ be the set of all such intervals. For every time t , define

$$\text{ACTIVE}(t) := \{I \in \mathcal{I} : t \in I\}$$

as the set of intervals in $\mathcal{I}$ that contain t . It is immediate that $|\text{ACTIVE}(t)| = O(\log t)$. A depiction of the set $\mathcal{I}$ is shown in Figure 1.

For every interval $I = [p, q] \in \mathcal{I}$, we keep an instance of FTL F_I . Let

$$W_t(k, I) = \sum_{s=p}^t r_s(k) - r_s(x_s)$$

Figure 1: Depiction of interval sets $\mathcal{I}_n$

denote the weight F_I assigns to expert k at time t , where x_t is the expert, or version, chosen by our algorithm. Notice that this differs slightly from the traditional FTL definition as we subtract the reward of the action taken by our algorithm. In doing so, we ensure that the scale of the weights $W_t(k, I)$ is roughly the same for all intervals I , regardless of their length.

The expert and interval chosen by FTBI is given by

$$(x_t, I_t) = \underset{k \in \{1,2\}, I \in \text{ACTIVE}(t)}{\operatorname{argmax}} W_{t-1}(k, I).$$

A full description of FTBI can be found in Algorithm 1. The intuition for our algorithm is the following: consider a segment $\mathcal{T}_i$ where rewards are i.i.d.; due to the properties of the FTL algorithm, any FTL instance beginning inside this segment will, in constant time, choose the optimal action k_i^* . Therefore, as long as our algorithm chooses FTL instances initialized inside a segment $\mathcal{T}_i$, FTBI will choose the optimal action. We will show that the choice of weights $W_t(k, I)$ will guarantee that only instances of FTL that observe i.i.d rewards will have large weights.

Algorithm 1 Follow-The-Best-Interval

Initialize $W_0(k, I) = 0$, for all $I \in \mathcal{I}$, $k \in \{1, 2\}$.
for $t = 1, 2, \dots, T$ **do**
 Choose $(x_t, I_t) = \operatorname{argmax}_{k \in \{1,2\}, I \in \text{ACTIVE}(t)} W_{t-1}(k, I)$
 Observe reward vector r_t
 Update $W_t(k, I) = W_{t-1}(k, I) + r_t(k) - r_t(x_t)$, for all $I \in \text{ACTIVE}(t)$, $k \in \{1, 2\}$.
end for

Remark. Since our algorithm is motivated by the A/B test problem, we focus on the case with two experts. However, the FTBI algorithm naturally extends to the setting with $K > 2$ experts. In that case, each interval needs to maintain K weights, $W_t(k, I)$, $k = 1, 2, \dots, K$. The algorithm chooses the expert and interval using:

$$(x_t, I_t) = \underset{k \in [K], I \in \text{ACTIVE}(t)}{\operatorname{argmax}} W_{t-1}(k, I).$$

4 ANALYSIS

In this section we provide regret guarantees for our algorithm. The main result of this paper is the following theorem.

Theorem 1. *There exists a universal constant C such that the shifting pseudo-regret of FTBI is bounded by*

$$R_T \leq \frac{C}{\Delta^3} N^2 \log^3 T.$$

As we can see, the regret bound of the FTBI is weaker than that of ANH, i.e., $O(\frac{N}{\Delta} \log T)$. However, as we will see in Section 5, FTBI almost always outperforms ANH in practice. We believe this is an indication that our upper bounds are not tight. It remains an open problem to provide a lower bound on the performance of our algorithm.

Figure 2: Depiction of the cover. We show the covering of the interval $[1, 30]$ by intervals $[1, 1], [2, 3], [4, 7], [8, 15], [16, 23], [24, 27], [28, 29], [30, 30] \in \mathcal{I}$.

4.1 Proof Sketch

We begin by defining a disjoint cover of each segment $\mathcal{T}_i$ by elements of $\mathcal{I}$. This cover was first introduced by [8] to analyze their strongly adaptive algorithm.

Theorem 2 ([8], Geometric Covering). *Any interval $I \subseteq [1, T]$ can be partitioned into two finite sequences of disjoint and consecutive intervals, denoted by $I^{-m}, \dots, I^0 \in \mathcal{I}$ and $I^1, \dots, I^h \in \mathcal{I}$ such that*

$$(\forall i \geq 1), |I^{-i}|/|I^{-i+1}| \leq 1/2,$$

$$(\forall i \geq 2), |I^i|/|I^{i-1}| \leq 1/2.$$

The above theorem shows that any interval I in $[1, T]$ can be partitioned by at most $2 \log_2 |I|$ intervals from $\mathcal{I}$. Since the first sequence $I^{-m}, \dots, I^0$ has geometrically increasing lengths and the second sequence $I^1, \dots, I^h$ has geometrically decreasing lengths, we call this partitioning the *geometric covering* of interval I . A depiction of this cover can be seen in Figure 2.

Denote by $C \subseteq \mathcal{I}$ the collection of the geometric coverings for all segments $\mathcal{T}_1, \dots, \mathcal{T}_N$. We know that, C forms a partition of $[1, T]$, and there are at most $2N \log_2 T$ intervals in C . Furthermore, since each interval $I \in C$ is completely contained in $\mathcal{T}_j$ for some j , the rewards are i.i.d. for any interval $I \in C$. The main idea behind our proof is to analyze the regret of each interval $I \in C$, and to sum them up to obtain the overall regret of FTBI. At a first glance bounding the regret for $I \in C$ is straightforward since the observed rewards are i.i.d. However, the FTL instance F_I over interval I is not isolated. In particular, it could *compete* with other instances F_J over intervals $J \supset I$. Therefore we must show that no FTL instance accumulates a large weight in which case F_I can *catch up* to the weights of F_J in a small number of rounds.

For any $I = [p, q] \in C$, let $\mathcal{L}_I = \{J \in \mathcal{I} : J \supseteq I\}$. That is, $\mathcal{L}_I$ consists of all intervals in $\mathcal{I}$ that contain I . One can verify that, from the definition of $\mathcal{I}$, if an interval $J = [p', q'] \in \mathcal{I}$ overlaps with I , i.e., $I \cap J \neq \emptyset$, and it was active before I , i.e., $p' < p$, then it must be that $I \subset J$, i.e., $J \in \mathcal{L}_I$. In view of this, it is the weight of the intervals in $\mathcal{L}_I$ that we need to bound. Let

$$H_I^k = \operatorname{argmax}_{J \in \mathcal{L}_I} W_{p-1}(k, J)$$

denote the interval in $\mathcal{L}_I$ that assigns the highest weight to action k at the beginning of interval I . Let $\mathcal{F}_t$ denote the sigma algebra generated by all history up to time t and denote by R_I the pseudo-regret of FTBI over interval I conditioned on all the previous history, i.e.,

$$R_I := \max_{k \in \{1,2\}} \mathbb{E} \left[\sum_{t=p}^q r_t(k) - r_t(x_t) \mid \mathcal{F}_{p-1} \right].$$

Theorem 3. For any $I = [p, q] \in C$, we have

$$R_I \leq \mathbb{E}[\max\{W_q(1, H_I^1), W_q(2, H_I^2), 0\} \mid \mathcal{F}_{p-1}] \\ \leq \max\{W_{p-1}(1, H_I^1), W_{p-1}(2, H_I^2), 0\} + \frac{28}{\Delta^3} \log T + O(1). \quad (1)$$

We defer the full proof to the Appendix. This theorem shows that the conditional regret in each interval depends on the maximum weight accumulated over all FTL instances before time p plus $\frac{28}{\Delta^3} \log T$. More importantly, inequality (1) shows that at the end of interval I the maximum weight over all active FTL instances and all actions can only increase by an additive constant.

We can now prove Theorem 1. First, by construction of the geometric covering, we have $m := |C| \leq 2N \log_2 T$. Let $I(1), \dots, I(m)$ be an ordered enumeration of the intervals in C . For every interval $I = [p, q] \in C$, let $W^*(I) = \max\{W_{p-1}(1, H_I^1), W_{p-1}(2, H_I^2), 0\}$. Applying inequality (1) recursively it is easy to see that:

$$\mathbb{E}[W^*(I(n))] \leq (n-1) \left[\frac{28}{\Delta^3} \log T + O(1) \right],$$

which again in view of (1) yields

$$\mathbb{E}[R_{I(n)}] \leq n \left[\frac{28}{\Delta^3} \log T + O(1) \right].$$

Thus

$$R_T = \sum_{n=1}^m \mathbb{E}[R_{I(n)}] \leq \frac{1}{2} m(m+1) \left[\frac{28}{\Delta^3} \log T + O(1) \right].$$

The proof is completed using the fact that $m \leq 2N \log_2 T$.

5 EXPERIMENTS

We conduct extensive simulations comparing our algorithm (FTBI) to AdaNormalHedge (ANH) on both synthetic and real world datasets. In order to obtain an efficient implementation of ANH, we use the streaming trick mentioned in [11, 19].

5.1 Synthetic Data

The purpose of these simulations is to understand the advantages of FTBI over ANH in a controlled environment for different parameters representing the reward gap Δ , number of shifts N , and round length $|\mathcal{T}_i|$. For each experiment we calculate the regret: $\bar{R}_T = \sum_{i=1}^N \sum_{t=\tau_{i-1}+1}^{\tau_i} r_t(k_i^*) - r_t(x_t)$. Each experiment is replicated 20 times and we report the average result. The error bars represent one standard deviation. For every experiment the rewards of the two experts over the i -th segment $\mathcal{T}_i$ are sampled i.i.d. from a Bernoulli distribution with parameters $p_{i,1}, p_{i,2}$.

We explore the performance of FTBI as a function of the reward gap, Δ , the number of shifts, N , round length, T_i , and the number of experts, K .

5.1.1 Sensitivity to the reward gap Δ . For this experiment, we vary the reward gap over the set $\Delta \in \{0.06, 0.17, 0.28, 0.39, 0.5\}$ and choose parameters $p_{i,k}$ such that $|p_{i,1} - p_{i,2}| = \Delta$. We allow for two intervals $\mathcal{T}_i$ with 800 rounds per interval, and plot the results in Figure 3(a). Observe that FTBI consistently outperforms ANH, and the regret exhibits a significantly better dependence on Δ than guaranteed by Theorem 1.

5.1.2 Sensitivity to the number of shifts, N . We vary the number of shifts N from 2 to 32. For each setting, we sample parameters $p_{i,1}, p_{i,2}$ from a uniform distribution in $[0, 1]$, and fix the number of rounds per interval to 800. Observe that a new draw of p_i 's initiates a shift with high probability as $p_{i,k} \neq p_{i+1,k}$. We compare the

Figure 3: Performance on synthetic data. Regret as a function of (a) reward gap Δ , (b) number of shifts N , (c) round length $|\mathcal{T}_i|$, (d) number of experts K .

Figure 4: Performance on synthetic data. Regret as a function of number of shifts N and reward gap Δ , with N and Δ chosen as functions of T . (a) Fix $\Delta = 0.1$ and let $N = \sqrt{T}$. (b) Fix $N = 10$ and let $\Delta = 1/\sqrt{T}$.

regret of FTBI and ANH in Figure 3(b). Notice that although our theoretical bound shows a quadratic dependence on the number of shifts, we do not see this unfavorable bound empirically, and again FTBI consistently outperforms ANH.

5.1.3 Sensitivity to the round length, T_i . We fix the number of intervals to 2. On the first interval expert 1 has mean reward 0.7 and expert 2 has reward 0.5. The means are swapped in the second interval. We vary the length of each interval over the set $\{10^1, 10^2, 10^3, 10^4\}$, and show the results in Figure 3(c). This experiment captures the main downside of ANH. Specifically, ANH is slow to adjust to a shift in distributions. Indeed, the longer one arm dominates the other, the longer it takes ANH to switch when a distribution shift occurs. By contrast, FTBI adjusts much faster.

5.1.4 Sensitivity to the number of experts, K . As mentioned in Section 3, although our main focus is the A/B test problem, as an expert learning algorithm, FTBI also applies to scenarios where there are $K > 2$ experts. In this experiment, we study the performance of FTBI when we vary the number of experts K . We sample the parameters $p_{i,k}$ ($k = 1, 2, \dots, K$) i.i.d. from a uniform distribution in $[0, 1]$.

We fix the length of each segment $\mathcal{T}_i$ to 800 rounds, and vary the number of experts k from 2 to 32. We show the results in Figure 3(d). Although FTBI does not have theoretical guarantees in this setting, it consistently outperforms ANH. In addition, FTBI appears to scale better than ANH as the number of experts grows.

5.1.5 Sensitivity to the number of shifts and reward gap as functions of T . We also test the performance of FTBI and ANH in the settings where the number of shifts N and/or the reward gap Δ are functions of T . These settings simulate practical scenarios where there are relatively large number of shifts and) small reward gap between the optimal and sub-optimal experts. More importantly, it is in these scenarios that our theoretical bounds become trivial as setting $N = \sqrt{T}$ or $\Delta = \frac{1}{\sqrt{T}}$ makes the bound of Theorem 1 simply $O(T)$. In theory our algorithm should perform worse than ANH. However, as seen in Figure 4 this is not the case. For Figure 4(a), we fixed $\Delta = 0.1$ and let $N = \sqrt{T}$. We observe that FTBI achieves significantly smaller regret than ANH. Again, this is due to the fact that larger times without shifts are detrimental to the performance of ANH. For Figure 4(b), we fix $N = 10$ and let $\Delta = 1/\sqrt{T}$ since this setup corresponds to experts with similar performance both algorithms achieve similar regret.

5.2 Evaluation on Public Datasets

We now evaluate our algorithm by simulating the outcome of an A/B test based on public time series data. The methodology is the following: given a time series we consider two prediction systems. The first one predicts that the value will be above a fixed threshold τ and the other one predicts the it will be below τ . The reward of each system is 1 when the system is correct and 0 otherwise. The goal is to find a combination of systems that yields the largest cumulative reward. As mentioned in the introduction, this setup matches that of expert learning with two experts.

The first time series we consider is the air quality data set [17] (<https://archive.ics.uci.edu/ml/datasets/Beijing+PM2.5+Data>) consisting of hourly measurements of air quality in Beijing. The threshold of our systems is given by $\tau = 75$ corresponding to the rounded median of the air quality index. The second data set consists of measurements of electric power usage (<https://archive.ics.uci.edu/ml/datasets/Individual+household+electric+power+consumption>) and the threshold is given by $\tau = 0.5$.

We compare the performance of FTBI, ANH, and FTL algorithms in these tasks. Before presenting the results of these comparisons we first show why these data sets benefit from the adaptive nature of FTBI. In Figure 5(a) we plot the difference of the cumulative reward of the thresholding systems. Notice that if one system consistently dominated the other (a favorable case for FTL) we would see a monotonically increasing curve. Instead, we observe an oscillatory behavior, which indicates that the best system changes several times over the time period.

We run the three algorithms in these data sets and show the total reward of FTBI and ANH as a relative improvement over FTL in Figure 5(b). There, we see that both ANH and FTL largely outperform FTL with improvements of up to 85%, which validates the fact that these algorithms are able to adapt to the changes in the best system. Furthermore, our algorithm outperforms ANH. In the first data set it is considerably better, while in the second

Figure 5: (a) Difference of cumulative rewards of both systems. (b) Total reward of FTBI and ANH as a relative improvement over the total reward of FTL.

one the difference is not as large. By looking at the corresponding cumulative reward plots we see that the first task is noisier than the second one which leads us to believe that FTBI performs better under noisy conditions.

5.3 Evaluation on an Advertising Exchange Platform

Finally, to showcase the advantage of our method in practice, we apply the FTBI algorithm to compare two pricing mechanisms on an advertising exchange platform using real auction data. We consider the problem of setting reserve (or minimum) prices in the auctions, with the goal of increasing revenue. Learning good reserve prices is a notoriously challenging problem, with a myriad of algorithms designed to solve this task [5, 7, 20, 21].

In this experiment, we collect a sample of auctions (highest and second highest bid) for an eight day window from four different publishers, see the total number of auctions in Table 1. We consider an A/B test over two algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$ for setting the reserve price; we treat them as black boxes for the purpose of the experiment. At every time t , we can compute the revenues¹ $r_t(1)$ and $r_t(2)$ that could be obtained by each of these algorithms. (We ignore the question of incentive compatibility of learning in auctions, as it is beyond the scope of this work.)

Publisher	Number of auctions	Publisher	Number of auctions
A	20,745	B	115,341
C	64,524	D	34,789

Table 1: Total number of auctions in our data set

Since the reserve price must be set prior to running the auction, our goal is to select one of these algorithms at each point of time t , to maximize the overall revenue. We run the Follow-The-Leader (FTL) algorithm, AdaNormalHedge (ANH) and Follow-The-Best-Interval (FTBI) over these sequences of rewards. We present the relative revenue lift of FTBI (defined as the relative difference in revenue) over the two other algorithms in Figure 6, where we treat each day as a separate experiment. The box limits represent the

¹The revenue in a particular auction equals the amount of money that the winner of the auction pays to the publisher.

Figure 6: Performance on Auction data. Revenue lift comparisons in percent improvement. (a) Lift of FTBI against ANH. (b) Lift of FTBI against FTL. The box limits represent the 25 and 75 percent quantiles while the whiskers represent the 10 and 90 percent quantiles.

Figure 7: Performance on Auction data. Difference in the cumulative revenue of algorithms $\mathcal{A}_1$ and $\mathcal{A}_2$.

25 and 75 percent quantiles while the whiskers represent the 10 and 90 percent quantiles. Observe from Figure 6(a) is that FTBI is consistently better than ANH, in some cases improving total revenue by more than 4%.

Figure 6(b) shows that FTBI does not always compare favorably with respect to FTL. This can be explained by the fact that in these sequences, one algorithm is consistently better than the other throughout the day. It is well known that FTL has the best regret guarantees in this, non-switching, scenario, and FTBI essentially matches its performance on publishers A and C. In cases where the optimum strategy switches between $\mathcal{A}_1$ and $\mathcal{A}_2$, FTBI tends to outperform FTL, sometimes significantly.

To investigate this further, in Figure 7, we plot the difference in the cumulative revenue of $\mathcal{A}_1$ and $\mathcal{A}_2$ as a function of time for each publisher for four different days. When the time series increases, it represents a segment of time where $\mathcal{A}_1$ was consistently better than $\mathcal{A}_2$; the opposite is true when the series decreases. In particular, notice that, for publishers A and C, $\mathcal{A}_1$ consistently outperforms $\mathcal{A}_2$, explaining the results in Figure 6(b). On the other hand, the best algorithm in publisher B switches considerably in days 2 and 4 and there is a clear switch for publisher D in day 2. It is in these cases that we obtain the most benefits from an adaptive algorithm such as FTBI.

6 CONCLUSION

In this work, we study the setting of monitoring non-stationary A/B tests, and formulate an expert learning framework for this problem. We develop the Follow-The-Best-Interval (FTBI) algorithm and prove theoretical bounds on its regret. Our empirical evaluation suggests that it achieves lower regret than the state-of-the-art algorithm across many synthetic and real world datasets. We suggest improving the theoretical guarantee of FTBI as a future research direction.

A PROOF OF THEOREM 3

To simplify notation we fix $I = [p, q] \in \mathcal{C}$ and let $H^k := H_I^k$ for $k = 1, 2$. We define $W_1 := W_{p-1}(1, H^1)$ and $W_2 := W_{p-1}(2, H^2)$, we also let $\mathbb{E}$ denote the conditional expectation with respect to $\mathcal{F}_{p-1}$, and $\mathbb{P}$ denote the conditional probability with respect to $\mathcal{F}_{p-1}$. Finally, we assume, without loss of generality that expert 1 is optimal over the segment I .

The proof of Theorem 3 depends on the following three technical lemmas.

Lemma 1. *If $W_1 \geq W_2$, we have*

$$R_I \leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1] \leq \frac{6}{\Delta^2} \log T + O(1). \quad (2)$$

Lemma 2. *If $W_2 > W_1$, we have*

$$R_I \leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1] \leq \left(1 + \frac{2}{\Delta}\right)(W_2 - W_1) + \frac{16}{\Delta^2} \log T + O(1). \quad (3)$$

Lemma 3. *If $W_2 - W_1 > \frac{6}{\Delta^2} \log T$, we have*

$$R_I \leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1] \leq W_2 - W_1 + \frac{6}{\Delta^2} \log T + O(1). \quad (4)$$

The proof of Theorem 3 is now immediate. Combining Lemma 2 and Lemma 3, we know that for $W_2 > W_1$, we always have

$$\begin{aligned} R_I &\leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1] \\ &\leq W_2 - W_1 + \left(\frac{12}{\Delta^3} + \frac{16}{\Delta^2}\right) \log T + O(1) \\ &\leq W_2 - W_1 + \frac{28}{\Delta^3} \log T + O(1). \end{aligned} \quad (5)$$

Moreover, by definition of H^1 and since $I \in \mathcal{L}_I$ we must have $W_1 \geq W_{p-1}(1, I) = 0$. Therefore by Lemma 1, Lemma 2 and the fact that W_1 is $\mathcal{F}_{p-1}$ measurable we have:

$$\begin{aligned} R_I &\leq R_I + W_1 \leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\}] \\ &\leq \max\{W_1, W_2, 0\} + \frac{28}{\Delta^3} \log T + O(1), \end{aligned}$$

where the last inequality follows from (2) and (3).

B PROOF OF LEMMA 1

By definition of regret we have

$$\begin{aligned} R_I &= \mathbb{E} \left[\sum_{t=p}^q r_t(1) - r_t(x_t) \right] = \mathbb{E}[W_q(1, H^1) - W_1] \\ &\leq \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1] := \tilde{R}_I. \end{aligned}$$

Fix $\delta > 0$, and let $\tau = \inf\{q > t \geq p : W_t(1, H^1) \geq W_1 + \delta\}$, and $\tau = \infty$ if $W_t(1, H^1) < W_1 + \delta$ for all $t < q$. Define the following two events:

$$X := \{W_q(1, H^1) \geq W_q(2, H^2)\},$$

$$Y := \{\tau < q \wedge x_t = 1 \forall t \in [\tau + 1, q]\} \cup \{\tau = \infty\}.$$

A simple observation is that, if Y happens, we must have $W_t(1, H^1) \leq W_1 + \delta + 1$ for all $t \in I$. Indeed, for $t < \tau$, $W_t(1, H^1) < W_1 + \delta$ by definition. Thus $W_\tau(1, H^1) \leq W_1 + \delta + 1$, and then $W_t(1, H^1)$ remains unchanged as $x_t = 1$ for all $t \in [\tau + 1, q]$.

Therefore, conditioned on the event $X \cap Y$, we always have $\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1 \leq \delta + 1$, and we can bound $\tilde{R}_I$ as

$$\begin{aligned} \tilde{R}_I &= \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1 \mid X \cap Y] \mathbb{P}\{X \cap Y\} \\ &\quad + \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1 \mid \bar{X} \cup \bar{Y}] \mathbb{P}\{\bar{X} \cup \bar{Y}\} \\ &\leq \delta + 1 + 2T \mathbb{P}\{\bar{X} \cup \bar{Y}\}. \end{aligned}$$

Moreover, since $W_1 \geq W_2$, we have

$$\mathbb{P}\{\bar{X}\} \leq \mathbb{P}\left\{\sum_{t=p}^q r_t(1) - r_t(2) < 0\right\} \leq e^{-\frac{1}{2}\Delta^2|I|},$$

where we have used the fact that $\mathbb{E}[r_t(1) - r_t(2)] \geq \Delta$ and Hoeffding's inequality.

Let us bound $\mathbb{P}\{\bar{Y}\}$. Notice that event Y does not happen if and only if $\tau < q$ and there exists $\tau' > \tau$ such that $x_{\tau'} = 2$ and $x_t = 1$ for $t \in [\tau + 1, \tau' - 1]$. Then $W_t(1, H^1)$ remains unchanged for $t = \tau, \dots, \tau' - 1$ and $W_t(1, H^1) \geq W_1 + \delta$ for every $t \in [\tau, \tau' - 1]$. Thus, the previous event can happen if and only if there exists $t \in [p, q]$ such that $W_{t-1}(1, H^1) \geq W_1 + \delta$ and $x_t = 2$. Therefore, let

$$E_t := \{W_{t-1}(1, H^1) \geq W_1 + \delta \wedge x_t = 2\},$$

for all $t \in [p + 1, q]$. It is easy to see that $\bar{Y} = \bigcup_{t=p+1}^q E_t$. Also, if $t < p + \delta$, then $W_{t-1}(H^1) < W_1 + \delta$. Therefore, we need only bound the probability of E_t for $t \geq p + \delta$. Recall in Algorithm 1, we define

$$(x_t, I_t) = \underset{k \in \{1, 2\}, \tilde{I} \in \text{ACTIVE}(t)}{\operatorname{argmax}} W_{t-1}(k, \tilde{I}). \quad (6)$$

Therefore, we have

$$\begin{aligned} \mathbb{P}\{E_t\} &= \mathbb{P}\{W_{t-1}(1, H^1) \geq W_1 + \delta, x_t = 2\} \\ &\leq \mathbb{P}\{I_t = H^2, x_t = 2\} \\ &\quad + \sum_{\substack{J \in \text{ACTIVE}(t) \\ p_J > p}} \mathbb{P}\{W_{t-1}(1, H^1) \geq W_1 + \delta, I_t = J, x_t = 2\} \\ &\leq \mathbb{P}\{W_{t-1}(1, H^1) < W_{t-1}(2, H^2)\} \\ &\quad + \sum_{\substack{J \in \text{ACTIVE}(t) \\ p_J > p}} \mathbb{P}\{W_{t-1}(1, J) < W_{t-1}(2, J), W_{t-1}(2, J) \geq \delta\}, \end{aligned}$$

where we denote by p_J the starting time of interval J . Applying Hoeffding's inequality, we can bound the first term by $e^{-\frac{(t-p)\Delta^2}{2}} \leq e^{-\frac{\delta\Delta^2}{2}}$. For the second term notice that if $t - p_J < \delta$ then the probability of the summand is 0. Thus we can restrict ourselves to

intervals for which $t - p_J \geq \delta$ and we can bound the second term by

$$\begin{aligned} \sum_{\substack{J \in \text{ACTIVE}(t) \\ p_J > p, t - p_J \geq \delta}} \mathbb{P}\{W_{t-1}(1, J) < W_{t-1}(2, J)\} &\leq \sum_{\substack{J \in \text{ACTIVE}(t) \\ p_J > p, t - p_J \geq \delta}} e^{-\frac{(t-p_J)\Delta^2}{2}} \\ &\leq T e^{-\frac{\delta\Delta^2}{2}}, \end{aligned}$$

where we again used Hoeffding's inequality and we bound the number of active intervals by T . Thus, by union bound, we get

$$\mathbb{P}\{\bar{X} \cup \bar{Y}\} \leq e^{-\frac{\Delta^2|I|}{2}} + T^2 e^{-\frac{\delta\Delta^2}{2}}.$$

Then, by choosing $\delta = \frac{6}{\Delta^2} \log T$, and considering the two cases with $|I| \leq \delta$ and $|I| > \delta$, we get the desired result.

C PROOF OF LEMMA 2

The proof follows the same line of reasoning as the previous lemma. Let $\tilde{R}_I$ be as in the proof of Lemma 1. Let $C > 0$ be a constant to be chosen later and define $W = \max\{W_2 - W_1, C \log T\}$. Let X, Y and E_t be the events defined in the proof of Lemma 1. The same argument as before shows that

$$\tilde{R}_I \leq \delta + 1 + 2T \mathbb{P}\{\bar{X} \cup \bar{Y}\}$$

We proceed to bound $\mathbb{P}\{\bar{X} \cup \bar{Y}\}$. Without loss of generality, we assume that $\Delta = \mathbb{E}[r_t(1) - r_t(2)]$. We first consider the case where $|I| \geq \frac{2}{\Delta} W$.

For event $\bar{X}$, we have

$$\begin{aligned} \mathbb{P}\{\bar{X}\} &= \mathbb{P}\{W_q(1, H^1) < W_q(2, H^2)\} \\ &= \mathbb{P}\left\{W_1 + \sum_{t=p}^q r_t(1) < W_2 + \sum_{t=p}^q r_t(2)\right\} \\ &= \mathbb{P}\left\{\left(\sum_{t=p}^q r_t(1) - r_t(2)\right) - |I|\Delta < W_2 - W_1 - |I|\Delta\right\}, \end{aligned}$$

and since $|I| \geq \frac{2}{\Delta} W \geq \frac{2}{\Delta} (W_2 - W_1)$, by Hoeffding's bound we have

$$\mathbb{P}\{\bar{X}\} \leq \mathbb{P}\left\{\left(\sum_{t=p}^q r_t(1) - r_t(2)\right) - |I|\Delta < -\frac{1}{2}|I|\Delta\right\} \leq e^{-\frac{\Delta^2|I|}{8}}.$$

To bound $\mathbb{P}\{\bar{Y}\}$ we again bound the probability of the events E_t for $t \geq p + \delta$. Using the exact same technique as in Lemma 1 and letting $\delta = \frac{2}{\Delta} W$ we get

$$\mathbb{P}\{E_t\} \leq e^{-\frac{\Delta^2\delta}{8}} + T e^{-\frac{\Delta^2\delta}{2}}.$$

Then, by union bound, we obtain

$$\mathbb{P}\{\bar{X} \cup \bar{Y}\} \leq e^{-\frac{\Delta^2|I|}{8}} + T e^{-\frac{\Delta^2\delta}{8}} + T^2 e^{-\frac{\Delta^2\delta}{2}}.$$

Recall that $\delta = \frac{2}{\Delta} W = \frac{2}{\Delta} \max\{W_2 - W_1, C \log T\}$ and $|I| \geq \frac{2}{\Delta} W = \delta$. A simple calculation verifies that setting $C = \frac{8}{\Delta}$, yields

$$\begin{aligned} \tilde{R}_I &\leq \delta + O(1) \leq \frac{2}{\Delta} (W_2 - W_1 + C \log T) + O(1) \\ &\leq \frac{2}{\Delta} (W_2 - W_1) + \frac{16}{\Delta^2} \log T + O(1). \end{aligned}$$

On the other hand if $|I| < \frac{2}{\Delta} W$, we have

$$W_q(1, H^1) - W_1 \leq \frac{2}{\Delta} W, \quad W_q(2, H^2) - W_1 \leq \frac{2}{\Delta} W + W_2 - W_1.$$

Therefore we must have

$$\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1 \leq \frac{2}{\Delta} W + W_2 - W_1,$$

using the definition of W and bounding the max by a sum yields

$$\tilde{R}_I \leq \left(1 + \frac{2}{\Delta}\right)(W_2 - W_1) + \frac{16}{\Delta^2} \log T + O(1).$$

D PROOF OF LEMMA 3

Using the previous notation, we strive to bound

$$\tilde{R}_I := \mathbb{E}[\max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1]$$

under the assumption that $W_2 - W_1 > C \log T$ for $C = \frac{6}{\Delta^2}$. Let $\xi = \max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_1$ and define the stopping time

$$\eta := \inf\{q \geq t \geq p \mid \exists J \in \text{ACTIVE}(t) \text{ with } W_t(1, J) > W_t(2, H^2)\},$$

with $\eta = \infty$ if the event does not occur. Notice that by definition of η , expert 1 is chosen for the first time at time $\eta + 1$. We consider the following three disjoint events:

$$\mathcal{U} := \{\eta = \infty\}$$

$$\mathcal{V} := \{\eta < \infty \wedge W_\eta(1, H^1) \leq W_\eta(2, H^2)\}$$

$$\mathcal{W} := \{\eta < \infty \wedge W_\eta(1, H^1) > W_\eta(2, H^2)\}.$$

We make the following observations:

- (i) $W_t(2, H^2) = W_2$ for $t \leq \eta$. Indeed, since by definition of η we have $x_t = 2$ for $t \leq \eta$ it follows that $W_t(2, H^2)$ remains constant.
- (ii) If $\mathcal{U}$ happens then $\xi = W_2 - W_1$. By definition of η we must have $W_q(1, H^1) \leq W_q(2, H^2) = W_2$, which implies the statement.
- (iii) If $\mathcal{W}$ happens then we must have

$$W_{\eta-1}(1, H^1) \leq W_{\eta-1}(2, H^2) = W_2,$$

and this easily implies

$$W_2 < W_\eta(1, H^1) \leq W_2 + 1.$$

- (iv) $\mathbb{P}[\eta = l] = 0$ if $l - (p - 1) \leq C \log T$. This is immediate since W_1 is the largest weight over expert 1 at time $p - 1$ and by assumption $W_2 - W_1 > C \log T$.

Using observation (ii) can write $\tilde{R}_I$ as follows:

$$\begin{aligned} \tilde{R}_I &= \mathbb{E}[\xi \mid \mathcal{U}] \mathbb{P}\{\mathcal{U}\} + \mathbb{E}[\xi \mid \mathcal{V}] \mathbb{P}\{\mathcal{V}\} + \mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}] \\ &= (W_2 - W_1) \mathbb{P}\{\mathcal{U}\} + \mathbb{E}[\xi \mid \mathcal{V}] \mathbb{P}\{\mathcal{V}\} + \mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}]. \end{aligned} \quad (7)$$

We now proceed to bound the last two summands in the above expression. The main intuition behind our proof is that since we have $W_2 - W_1 > C \log T$ it is unlikely that $W_\eta(1, J) > W_\eta(2, H^2)$ for any interval different from H^1 . This will imply $\mathbb{P}\{\mathcal{V}\}$ is small. On the other hand if $\mathcal{W}$ occurs then the algorithm behavior after η is similar to FTL and thus the weights $W_t(1, H^1)$ and $W_t(2, H^2)$ does not significantly increase for $t > \eta$.

We proceed to bound $\mathbb{P}\{\mathcal{V}\}$. Let $\bar{J} = [p_{\bar{J}}, q_{\bar{J}}] \in \mathcal{I}$ be such that $W_\eta(1, \bar{J}) > W_\eta(2, H^2)$. Using the same argument of observation (iv) we must have $|\bar{J}| \geq C \log T$. Furthermore, by definition of $\mathcal{V}$, we must have $p_{\bar{J}} > p$ (again, we use p_J to denote the starting time of interval $J \in \mathcal{I}$). Indeed, if $\bar{J}$ starts at or before p , then by definition of H^1 , $W_{p-1}(1, H^1) \geq W_{p-1}(1, \bar{J})$ which would imply $W_\eta(1, H^1) \geq W_\eta(1, \bar{J})$ contradicting the definition of $\mathcal{V}$. Finally,

by construction of $\mathcal{I}$ if $\bar{J}$ does not start at p then we must have $p_J - p \geq |\bar{J}| \geq C \log T$. Therefore we have:

$$\begin{aligned} \mathbb{P}\{\mathcal{V}\} &\leq \mathbb{P}\{W_\eta(1, H_1) < W_\eta(1, \bar{J})\} \\ &= \mathbb{P}\left\{W_1 + \sum_{t=p}^{\eta} r(1) - r(2) < \sum_{t=p_J}^{\eta} r(1) - r(2)\right\} \\ &= \mathbb{P}\left\{W_1 + \sum_{t=p}^{p_J-1} r(1) - r(2) < 0\right\} \\ &\leq \mathbb{P}\left\{\sum_{t=p}^{p_J-1} r(1) - r(2) < 0\right\}. \end{aligned}$$

We can use the union bound and the observation that $p_J - p \geq C \log T$ to bound the above probability as

$$\begin{aligned} \mathbb{P}\left\{\sum_{t=p}^{p_J-1} r(1) - r(2) < 0\right\} &\leq \sum_{\substack{J \in \mathcal{I} \\ p_J - p \geq C \log T}} \mathbb{P}\left\{\sum_{t=p}^{p_J-1} r(1) - r(2) < 0\right\} \\ &\leq 2Te^{-\frac{\Delta^2(p_J-p)}{2}}. \end{aligned}$$

where applied Hoeffding's inequality, the fact that $\mathbb{E}[r(1) - r(2)] \geq \Delta$ and the fact that $|\mathcal{I}| \leq 2T$. Finally, using the fact that $p_J - p \geq C \log T$ and replacing the value of C we see that

$$\mathbb{P}\{\mathcal{V}\} \leq \frac{2}{T^2} \quad (8)$$

We proceed to bound $\mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}]$. Notice that from observation (iii) we have

$$\begin{aligned} \xi &= \max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_2 + W_2 - W_1 \\ &\leq \max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_\eta(1, H^1) + 1 + W_2 - W_1. \end{aligned}$$

Let $\xi_\eta = \max\{W_q(1, H^1), W_q(2, H^2), 0\} - W_\eta(1, H^1)$. Fix $\delta > 0$ and notice that if $q - \eta \geq \delta$ we necessarily must have $\xi_\eta \leq \delta$. Therefore, we focus on the case where $q - \eta > \delta$. As in the two previous lemmas define $\tau = \inf\{q > t \geq \eta : W_t(1, H^1) \geq W_\eta(1, H^1) + \delta\}$ and $\tau = \infty$ if $W_t(1, H^1) < W_\eta(1, H^1) + \delta$ for all $t \in [p, q]$. Define also the events:

$$X := \{W_q(1, H^1) \geq W_q(2, H^2)\},$$

$$Y := \{\tau < q \wedge x_t = 1 \forall t \in [\tau + 1, q]\} \cup \{\tau = \infty\}.$$

With the same arguments as the previous two lemmas we can show that if $X \cap Y$ happens we must have $\xi_\eta < \delta + 1$. It follows that:

$$\begin{aligned} \mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}] &\leq (\delta + 2 + W_2 - W_1) \mathbb{P}\{\mathcal{W}\} \\ &\quad + \mathbb{E}[\xi_\eta + 1 + W_2 - W_1 \mid \bar{X} \cap \bar{Y}, \mathcal{W}] \mathbb{P}\{\bar{X} \cap \bar{Y}, \mathcal{W}\}. \end{aligned} \quad (9)$$

Notice that if $W_q(1, H_1) < W_q(2, H_2)$ and $\mathcal{W}$ happens we must have $\sum_{t=\eta+1}^q r_t(1) - r_t(2) < 0$. Using the tower property of conditional expectation we have

$$\begin{aligned} \mathbb{P}\{\bar{X}, \mathcal{W}\} &= \mathbb{E}[\mathbb{P}\{\bar{X}, \mathcal{W} \mid \eta\}] \leq \mathbb{E}\left[\mathbb{P}\left\{\sum_{t=\eta+1}^q r_t(1) - r_t(2) < 0 \mid \eta\right\}\right] \\ &\leq \mathbb{E}\left[e^{-\frac{\Delta^2(q-\eta)}{2}}\right] \leq e^{-\frac{\Delta^2 \delta}{2}}. \end{aligned}$$

Here for the last inequality we use Hoeffding's inequality and the fact that $q - \eta > \delta$ by assumption. Notice that, here, although we condition on η , i.e., the event $W_\eta(1, H^1) > W_\eta(2, H^2)$, this event only depends on the random rewards up to time η , and does not affect the independence of rewards starting from $\eta + 1$, and therefore we can apply Hoeffding's inequality. To bound the probability of $\bar{Y} \cup \mathcal{W}$ we again introduce the events

$$E_t := \{t > \eta, W_{t-1}(1, H^1) > W_\eta(1, H^1) + \delta \wedge x_t = 2\},$$

And notice that $\bar{Y} = \bigcup E_t$. Moreover notice that if $t - \eta \leq \delta$ then $\mathbb{P}\{E_t | \eta\} = 0$. Therefore we can assume that $t - \eta > \delta$. Define I_t as in (6). We can now bound $\mathbb{P}\{E_t, \mathcal{W}\}$ as follows:

$$\begin{aligned} & \mathbb{P}\{E_t, \mathcal{W}\} \\ &= \mathbb{E}[\mathbb{P}\{E_t, \mathcal{W} | \eta\}] \\ &= \mathbb{E}[\mathbb{P}\{W_{t-1}(1, H^1) > W_\eta(1, H^1) + \delta, x_t = 2, \mathcal{W} | \eta\}] \\ &\leq \mathbb{E}[\mathbb{P}\{W_{t-1}(1, H^1) < W_{t-1}(1, H^2), \mathcal{W} | \eta\}] \\ &\quad + \sum_{\substack{J \in \text{ACTIVE}(t) \\ p_J > p}} \mathbb{E}[\mathbb{P}\{W_{t-1}(1, H^1) > \delta, I_t = J, x_t = 2, \mathcal{W} | \eta\}]. \end{aligned} \quad (10)$$

Here we notice that since we condition on η , we have $W_\eta(1, H^1) > W_\eta(2, H^2)$. Therefore the even in the first summand of the above expression happens if and only if $\sum_{s=\eta+1}^{t-1} r_s(1) - r_s(2) \leq 0$. Using Hoeffding's bound we can thus bound the first term by

$$\mathbb{E}[\mathbb{P}\{W_{t-1}(1, H^1) < W_{t-1}(1, H^2), \mathcal{W} | \eta\}] \leq \mathbb{E}[e^{-\frac{\Delta^2(t-\eta)}{2}}] \leq e^{-\frac{\Delta^2 \delta}{2}}. \quad (11)$$

We proceed to bound the second term in (10). We split the sum to cases where $J = [p_J, q_J]$ is such that $p_J \leq \eta$ and that $p_J > \eta$. For the first case, notice that since $W_t(2, H^2)$ remains unchanged for $t = p, \dots, \eta$, $x_t = 2$ for all $t = p, \dots, \eta$, and $p_J > p$, we have $W_\eta(2, H^2) > W_\eta(2, J) = 0$. Therefore, for all $t > \eta$, $W_t(2, H^2) > W_t(2, J)$, and thus $I_t = J$ cannot happen. For the second case, if $p_J > \eta$ notice that the event can only happen if $W_{t-1}(2, J) > \delta$ and $W_{t-1}(2, J) > W_{t-1}(1, J)$. This implies that $t - p_J > \delta$ and the same argument as before shows

$$\mathbb{P}\{W_{t-1}(1, H^1) > \delta, I_t = J, x_t = 2, \mathcal{W} | \eta\} \leq e^{-\frac{\delta \Delta^2}{2}}. \quad (12)$$

Using inequalities (11) and (12) and the fact that there are at most T active intervals at time t we see that (10) can be bounded as:

$$\mathbb{P}\{E_t, \mathcal{W}\} \leq T e^{-\frac{\delta \Delta^2}{2}}.$$

By the union bound we thus have

$$P(\bar{X} \cup \bar{Y}, \mathcal{W}) \leq (1 + T^2) e^{-\frac{\Delta^2 \delta}{2}} \leq 2T^2 e^{-\frac{\Delta^2 \delta}{2}}.$$

Replacing this in (9) we see that:

$$\mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}] \leq (\delta + 2T^3 e^{-\frac{\Delta^2 \delta}{2}} + W_2 - W_1 + 2) \mathbb{P}\{\mathcal{W}\}$$

Finally, letting $\delta = \frac{6}{\Delta^2} \log T$ we obtain:

$$\mathbb{E}[\xi \mathbb{1}_{\mathcal{W}}] \leq \left(\frac{6}{\Delta^2} \log T + W_2 - W_1 + O(1) \right) \mathbb{P}\{\mathcal{W}\}. \quad (13)$$

Finally combining this bound with (7) and (8) we obtain the desired bound.

REFERENCES

- [1] Dmitry Adamskiy, Wouter M Koolen, Alexey Chernov, and Vladimir Vovk. 2012. A closer look at adaptive regret. In *International Conference on Algorithmic Learning Theory*. Springer, 290–304.
- [2] Dmitry Adamskiy, Manfred K Warmuth, and Wouter M Koolen. 2012. Putting Bayes to sleep. In *Advances in Neural Information Processing Systems*. 135–143.
- [3] Olivier Bousquet and Manfred K Warmuth. 2002. Tracking a small set of experts by mixing past posteriors. *Journal of Machine Learning Research* 3, Nov (2002), 363–396.
- [4] Nicolò Cesa-Bianchi, Yoav Freund, David Haussler, David P Helmbold, Robert E Schapire, and Manfred K Warmuth. 1997. How to use expert advice. *Journal of the ACM (JACM)* 44, 3 (1997), 427–485.
- [5] Nicolò Cesa-Bianchi, Pierre Gaillard, Claudio Gentile, and Sébastien Gerchinovitz. 2017. Algorithmic Chaining and the Role of Partial Feedback in Online Nonparametric Learning. In *Proceedings of COLT*. 465–481.
- [6] Nicolò Cesa-Bianchi, Pierre Gaillard, Gábor Lugosi, and Gilles Stoltz. 2012. Mirror descent meets fixed share (and feels no regret). In *Advances in Neural Information Processing Systems*. 980–988.
- [7] Nicolò Cesa-Bianchi, Claudio Gentile, and Yishay Mansour. 2015. Regret Minimization for Reserve Prices in Second-Price Auctions. *IEEE Trans. Information Theory* 61, 1 (2015), 549–564.
- [8] Amit Daniely, Alon Gonen, and Shai Shalev-Shwartz. 2015. Strongly Adaptive Online Learning. In *ICML*. 1405–1411.
- [9] Yoav Freund and Robert E Schapire. 1995. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*. Springer, 23–37.
- [10] Yoav Freund, Robert E Schapire, Yoram Singer, and Manfred K Warmuth. 1997. Using and combining predictors that specialize. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*. ACM, 334–343.
- [11] Elad Hazan and C Seshadhri. 2007. Adaptive algorithms for online decision problems. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 14.
- [12] Mark Herbster and Manfred K Warmuth. 1998. Tracking the best expert. *Machine learning* 32, 2 (1998), 151–178.
- [13] Mark Herbster and Manfred K Warmuth. 2001. Tracking the best linear predictor. *Journal of Machine Learning Research* 1, Sep (2001), 281–309.
- [14] Marcus Hutter and Jan Poland. 2005. Adaptive online prediction by following the perturbed leader. *Journal of Machine Learning Research* 6, Apr (2005), 639–660.
- [15] Kwang-Sung Jun, Francesco Orabona, Rebecca Willett, and Stephen Wright. 2016. Improved Strongly Adaptive Online Learning using Coin Betting. *arXiv preprint arXiv:1610.04578* (2016).
- [16] Adam Kalai and Santosh Vempala. 2005. Efficient algorithms for online decision problems. *J. Comput. System Sci.* 71, 3 (2005), 291–307.
- [17] Xuan Liang, Tao Zou, Bin Guo, Shuo Li, Haozhe Zhang, Shuyi Zhang, Hui Huang, and Song Xi Chen. 2015. Assessing Beijing's PM2.5 pollution: severity, weather impact, APEC and winter heating. *Proceedings of the Royal Society A* 471, 2182 (2015), 20150257.
- [18] Nick Littlestone and Manfred K Warmuth. 1994. The weighted majority algorithm. *Information and computation* 108, 2 (1994), 212–261.
- [19] Haipeng Luo and Robert E Schapire. 2015. Achieving All with No Parameters: AdaNormalHedge. In *COLT*. 1286–1304.
- [20] Andres Muñoz Medina and Sergei Vassilvitskii. 2017. Revenue Optimization with Approximate Bid Predictions. In *Proceedings of NIPS*. 1856–1864.
- [21] Mehryar Mohri and Andres Muñoz Medina. 2014. Learning Theory and Algorithms for revenue optimization in second price auctions with reserve. In *Proceedings of ICML*. 262–270.
- [22] Amir Sani, Gergely Neu, and Alessandro Lazaric. 2014. Exploiting easy data in online optimization. In *Advances in Neural Information Processing Systems*. 810–818.
- [23] Suhrid Satyal, Ingo Weber, Hye-young Paik, Claudio Di Ciccio, and Jan Mendling. 2017. AB-BPM: performance-driven instance routing for business process improvement. In *International Conference on Business Process Management*. Springer, 113–129.
- [24] Suhrid Satyal, Ingo Weber, Hye-young Paik, Claudio Di Ciccio, and Jan Mendling. 2018. AB Testing for Process Versions with Contextual Multi-armed Bandit Algorithms. In *International Conference on Advanced Information Systems Engineering*. Springer, 19–34.
- [25] Vladimir G Vovk. 1995. A game of prediction with expert advice. In *Proceedings of the eighth annual conference on Computational learning theory*. ACM, 51–60.
- [26] Manfred K Warmuth and Wouter M Koolen. 2014. Open Problem: Shifting Experts on Easy Data. In *COLT*. 1295–1298.
- [27] Lijun Zhang, Tianbao Yang, Rong Jin, and Zhi-Hua Zhou. 2017. Strongly Adaptive Regret Implies Optimally Dynamic Regret. *arXiv preprint arXiv:1701.07570* (2017).