
Contextual Bandit Learning with Predictable Rewards

Alekh Agarwal

alekh@cs.berkeley.edu

Miroslav Dudík

mdudik@yahoo-inc.com

Satyen Kale

sckale@us.ibm.com

John Langford

jl@yahoo-inc.com

Robert E. Schapire

schapire@cs.princeton.edu

Abstract

Contextual bandit learning is a reinforcement learning problem where the learner repeatedly receives a set of features (context), takes an action and receives a reward based on the action and context. We consider this problem under a realizability assumption: there exists a function in a (known) function class, always capable of predicting the expected reward, given the action and context. Under this assumption, we show three things. We present a new algorithm—Regressor Elimination—with a regret similar to the agnostic setting (i.e. in the absence of realizability assumption). We prove a new lower bound showing no algorithm can achieve superior performance in the worst case even with the realizability assumption. However, we do show that for *any* set of policies (mapping contexts to actions), there is a distribution over rewards (given context) such that our new algorithm has *constant* regret unlike the previous approaches.

1 Introduction

We are interested in the online contextual bandit setting, where on each round we first see a context $x \in \mathcal{X}$, based on which we choose an action $a \in \mathcal{A}$, and then observe a reward r . This formalizes several natural scenarios. For example, a common task at major internet engines is to display the best ad from a pool of options given some context such as information about the user, the page visited, the search query issued etc. The action set consists of the candidate ads and the reward is typically binary based on whether the user clicked the displayed ad or not. Another

natural application is the design of clinical trials in the medical domain. In this case, the actions are the treatment options being compared, the context is the patient’s medical record and reward is based on whether the recommended treatment is a success or not.

Our goal in this setting is to compete with a particular set of *policies*, which are deterministic rules specifying which action to choose in each context. We note that this setting includes as special cases the classical K -armed bandit problem (Lai and Robbins, 1985) and associative reinforcement learning with linear reward functions (Auer, 2003; Chu et al., 2011).

The performance of algorithms in this setting is typically measured by the *regret*, which is the difference between the cumulative reward of the best policy and the algorithm. For the setting with an arbitrary set of policies, the achieved regret guarantee is $\mathcal{O}(\sqrt{KT \ln(N/\delta)})$ where K is the number of actions, T is the number of rounds, N is the number of policies and δ is the probability of failing to achieve the regret (Beygelzimer et al., 2010; Dudík et al., 2011). While this bound has a desirably small dependence on the parameters T, N , the scaling with respect to K is often too big to be meaningful. For instance, the number of ads under consideration can be huge, and a rapid scaling with the number of alternatives in a clinical trial is clearly undesirable. Unfortunately, the dependence on K is unavoidable as proved by existing lower bounds (Auer et al., 2003).

Large literature on “linear bandits” manages to avoid this dependence on K by making additional assumptions. For example, Auer (2003) and Chu et al. (2011) consider the setting where the context x consists of feature vectors $x_a \in \mathbb{R}^d$ describing each action, and the expected reward function (given a context x and action a) has the form $w^T x_a$ for some fixed vector $w \in \mathbb{R}^d$. Dani et al. (2008) consider a continuous action space with $a \in \mathbb{R}^d$, without contexts, with a linear expected reward $w^T a$, which is generalized by Filippi et al. (2010) to $\sigma(w^T a)$ with a known Lipschitz-continuous link function σ . A striking aspect of the linear and generalized linear setting is that while the regret grows

Appearing in Proceedings of the 15th International Conference on Artificial Intelligence and Statistics (AISTATS) 2012, La Palma, Canary Islands. Volume XX of JMLR: W&CP XX. Copyright 2012 by the authors.

rapidly with the dimension d , it grows either only gently with the number of actions K (poly-logarithmic for Auer, 2003), or is independent of K (Dani et al., 2008; Filippi et al., 2010). In this paper, we investigate whether a weaker dependence on the number of actions is possible in more general settings. Specifically, we omit the linearity assumption while keeping the “realizability”—i.e., we still assume that the expected reward can be perfectly modeled, but do not require this to be a linear or a generalized linear model.

We consider an arbitrary class F of functions $f : (\mathcal{X}, \mathcal{A}) \rightarrow [0, 1]$ that map a context and an action to a real number. We interpret $f(x, a)$ as a predicted expected reward of the action a on context x and refer to functions in F as *regressors*. For example, in display advertising, the context is a vector of features derived from the text and metadata of the webpage and information about the user. The action corresponds to the ad, also described by a set of features. Additional features might be used to model interaction between the ad and the context. A typical regressor for this problem is a generalized linear model with a logistic link, modeling the probability of a click.

The set of regressors F induces a natural set of policies Π_F containing maps $\pi_f : \mathcal{X} \rightarrow \mathcal{A}$ defined as $\pi_f(x) = \operatorname{argmax}_a f(x, a)$. We make the assumption that the expected reward for a context x and action a equals $f^*(x, a)$ for some unknown function $f^* \in F$. The question we address in this paper is: Does this realizability assumption allow us to learn faster?

We show that for an arbitrary function class, the answer to the above question is “no”. The \sqrt{K} dependence in regret is in general unavoidable even with the realizability assumption. Thus, the structure of linearity or controlled non-linearity was quite important in the past works.

Given this answer, a natural question is whether it is at least possible to do better in various special cases. To answer this, we create a new natural algorithm, Regressor Elimination (RE), which takes advantage of realizability. Structurally, the algorithm is similar to Policy Elimination (PE) of Dudík et al. (2011), designed for the agnostic case (i.e., the general case without realizability assumption). While PE proceeds by eliminating poorly performing policies, RE proceeds by eliminating poorly predicting regressors. However, realizability assumption allows much more aggressive elimination strategy, different from the strategy used in PE. The analysis of this elimination strategy is the key technical contribution of this paper.

The general regret guarantee for Regressor Elimination is $\mathcal{O}(\sqrt{KT \ln(NT/\delta)})$, similar to the agnostic case. However, we also show that for *all* sets of policies Π there exists a set of regressors F such that $\Pi = \Pi_F$ and the regret of Regressor Elimination is $\mathcal{O}(\ln(N/\delta))$, i.e., independent of the number of rounds and actions. At the first sight, this seems to contradict our worst-case lower bound. This ap-

parent paradox is due to the fact that the same set of policies can be generated by two very different sets of regressors. Some regressor sets allow better discrimination of the true reward function, whereas some regressor sets will lead to the worst-case guarantee.

The remainder of the paper is organized as follows. In the next section we formalize our setting and assumptions. Section 3 provides our algorithm which is analyzed in Section 4. In Section 5 we present the worst-case lower bound, and in Section 6, we show an improved dependence on K in favorable cases. Our algorithm assumes the exact knowledge of the distribution over contexts (but not over rewards). In Section 7 we sketch how this assumption can be removed. Another major assumption is the finiteness of the set of regressors F . This assumption is more difficult to remove, as we discuss in Section 8.

2 Problem Setup

We assume that the interaction between the learner and nature happens over T rounds. At each round t , nature picks a context $x_t \in \mathcal{X}$ and a reward function $r_t : \mathcal{A} \rightarrow [0, 1]$ sampled i.i.d. in each round, according to a fixed distribution $D(x, r)$. We assume that $D(x)$ is known (this assumption is removed in Section 7), but $D(r|x)$ is unknown. The learner observes x_t , picks an action $a_t \in \mathcal{A}$, and observes the reward for the action $r_t(a_t)$. We are given a function class $F : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$ with $|F| = N$, where $|F|$ is the cardinality of F . We assume that F contains a perfect predictor of the expected reward:

Assumption 1 (Realizability). *There exists a function $f^* \in F$ such that $\mathbf{E}_{r|x}[r(a)] = f^*(x, a)$ for all $x \in \mathcal{X}$, $a \in \mathcal{A}$.*

We recall as before that the regressor class F induces the policy class Π_F containing maps $\pi_f : \mathcal{X} \rightarrow \mathcal{A}$ defined by $f \in F$ as $\pi_f(x) = \operatorname{argmax}_a f(x, a)$. The performance of an algorithm is measured by its *expected regret* relative to the best fixed policy:

$$\operatorname{regret}_T = \sup_{\pi_f \in \Pi_F} \sum_{t=1}^T [f^*(x_t, \pi_f(x_t)) - f^*(x_t, a_t)] .$$

By definition of π_f , this is equivalent to

$$\operatorname{regret}_T = \sum_{t=1}^T [f^*(x_t, \pi_{f^*}(x_t)) - f^*(x_t, a_t)] .$$

3 Algorithm

Our algorithm, Regressor Elimination, maintains a set of regressors that accurately predict the observed rewards. In each round, it chooses an action that sufficiently explores among the actions represented in the current set of regressors (Steps 1–2). After observing the reward (Step 3), the inaccurate regressors are eliminated (Step 4).

Sufficient exploration is achieved by solving the convex optimization problem in Step 1. We construct a distribution P_t over current regressors, and then act by first sampling a regressor $f \sim P_t$ and then choosing an action according to π_f . Similarly to the Policy Elimination algorithm of Dudík et al. (2011), we seek a distribution P_t such that the inverse probability of choosing an action that agrees with *any* policy in the current set is in expectation bounded from above. Informally, this guarantees that actions of any of the current policies are chosen with sufficient probabilities. Using this construction we relate the accuracy of regressors to the regret of the algorithm (Lemma 4.3).

A priori, it is not clear whether the constraint (3.1) is even feasible. We prove feasibility by a similar argument as in Dudík et al. (2011) (see Lemma A.1 in Appendix A). Compared with Dudík et al. (2011) we are able to obtain tighter constraints by doing a more careful analysis.

Our elimination step (Step 4) is significantly tighter than a similar step in Dudík et al. (2011): we eliminate regressors according to a very strict $\mathcal{O}(1/t)$ bound on the suboptimality of the least squares error. Under the realizability assumption, this stringent constraint will not discard the optimal regressor accidentally, as we show in the next section. This is the key novel technical contribution of this work.

Replacing $D(x)$ in the Regressor Elimination algorithm with the empirical distribution over observed contexts is straightforward, as was done in Dudík et al. (2011), and is discussed further in Section 7.

4 Regret Analysis

Here we prove an upper bound on the regret of Regressor Elimination. The proved bound is no better than the one for existing agnostic algorithms. This is necessary, as we will see in Section 5, where we prove a matching lower bound.

Theorem 4.1. *For all sets of regressors F with $|F| = N$ and all distributions $D(x, r)$, with probability $1 - \delta$, the regret of Regressor Elimination is $\mathcal{O}(\sqrt{KT \ln(NT/\delta)})$.*

Proof. By Lemma 4.1 (proved below), in round t if we sample an action by sampling f from P_t and choosing $\pi_f(x_t)$, then the expected regret is $\mathcal{O}(\sqrt{K \ln(NT/\delta)}/t)$ with probability at least $1 - \delta/2t^2$. The excess regret for sampling a uniform random action is at most $\mu \leq \frac{1}{\sqrt{T}}$ per round. Summing up over all the T rounds and taking a union bound, the total expected regret is $\mathcal{O}(\sqrt{KT \ln(NT/\delta)})$ with probability at least $1 - \delta$. Further, the net regret is a martingale; hence the Azuma-Hoeffding inequality with range $[0, 1]$ applies. So with probability at least $1 - \delta$ we have a regret of $\mathcal{O}(\sqrt{KT \ln(NT/\delta)} + \sqrt{T \ln(1/\delta)}) = \mathcal{O}(\sqrt{KT \ln(NT/\delta)})$. \square

Algorithm 1 Regressor Elimination

Input:

a set of reward predictors $F = \{f : (\mathcal{X}, \mathcal{A}) \rightarrow [0, 1]\}$
 distribution D over contexts, confidence parameter δ .

Notation:

$$\pi_f(x) := \operatorname{argmax}_{a'} f(x, a').$$

$$\hat{R}_t(f) := \frac{1}{t} \sum_{t'=1}^t (f(x_{t'}, a_{t'}) - r_{t'}(a_{t'}))^2.$$

For $F' \subseteq F$, define

$$A(F', x) := \{a \in \mathcal{A} : \pi_f(x) = a \text{ for some } f \in F'\}$$

$$\mu := \min\{1/2K, 1/\sqrt{T}\}.$$

For a distribution P on $F' \subseteq F$, define conditional distribution $P'(\cdot|x)$ on \mathcal{A} as:

w.p. $(1 - \mu)$, sample $f \sim P$ and return $\pi_f(x)$, and

w.p. μ , return a uniform random $a \in A(F', x)$.

$$\delta_t = \delta/2Nt^3 \log_2(t), \text{ for } t = 1, 2, \dots, T.$$

Algorithm:

$$F_0 \leftarrow F$$

For $t = 1, 2, \dots, T$:

1. Find distribution P_t on F_{t-1} such that

$$\forall f \in F_{t-1} : \mathbf{E}_x \left[\frac{1}{P_t'(\pi_f(x)|x)} \right] \leq \mathbf{E}_x [|A(F_{t-1}, x)|] \quad (3.1)$$

2. Observe x_t and sample action a_t from $P_t'(\cdot|x_t)$.
3. Observe $r_t(a_t)$.
4. Set

$$F_t = \left\{ f \in F_{t-1} : \hat{R}_t(f) < \min_{f' \in F_{t-1}} \hat{R}_t(f') + \frac{18 \ln(1/\delta_t)}{t} \right\}$$

Lemma 4.1. *With probability at least $1 - \delta_t N t \log_2(t) \geq 1 - \delta/2t^2$, we have:*

1. $f^* \in F_t$.
2. For any $f \in F_t$,

$$\mathbf{E}_{x,r} [r(\pi_f(x)) - r(\pi_{f^*}(x))] \leq \sqrt{\frac{200K \ln(1/\delta_t)}{t}}.$$

Proof. Fix an arbitrary function $f \in F$. For every round t , define the random variable

$$Y_t = (f(x_t, a_t) - r_t(a_t))^2 - (f^*(x_t, a_t) - r_t(a_t))^2.$$

Here, x_t is drawn from the unknown data distribution D , r_t is drawn from the reward distribution conditioned on x_t , and a_t is drawn from P_t' (which is defined conditioned on the choice of x_t and is independent of r_t). Note that this random variable is well-defined for *all* functions $f \in F$, not just the ones in F_t .

Let $\mathbf{E}_t[\cdot]$ and $\mathbf{Var}_t[\cdot]$ denote the expectation and variance conditioned on all the randomness up to round t . Using a form of Freedman's inequality from Bartlett et al. (2008) (see Lemma B.1) and noting that $Y_t \leq 1$, we get that with probability at least $1 - \delta_t \log_2(t)$, we have

$$\begin{aligned} & \sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}] - \sum_{t'=1}^t Y_{t'} \\ & \leq 4 \sqrt{\sum_{t'=1}^t \mathbf{Var}_{t'}[Y_{t'}] \ln(1/\delta_t) + 2 \ln(1/\delta_t)}. \end{aligned}$$

From Lemma 4.2, we see that $\mathbf{Var}_{t'}[Y_{t'}] \leq 4 \mathbf{E}_{t'}[Y_{t'}]$ so

$$\begin{aligned} & \sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}] - \sum_{t'=1}^t Y_{t'} \\ & \leq 8 \sqrt{\sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}] \ln(1/\delta_t) + 2 \ln(1/\delta_t)}. \end{aligned}$$

For notational convenience, define $X = \sqrt{\sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}]}$, $Z = \sum_{t'=1}^t Y_{t'}$, and $C = \sqrt{\ln(1/\delta_t)}$. The above inequality is equivalent to:

$$X^2 - Z \leq 8CX + 2C^2 \Leftrightarrow (X - 4C)^2 - Z \leq 18C^2.$$

This gives $-Z \leq 18C^2$. Since $Z = t(\hat{R}_t(f) - \hat{R}_t(f^*))$, we get that

$$\hat{R}_t(f^*) \leq \hat{R}_t(f) + \frac{18C^2}{t}.$$

By a union bound, with probability at least $1 - \delta_t N t \log_2(t)$, for all $f \in F$ and all rounds $t' \leq t$, we have

$$\hat{R}_{t'}(f^*) \leq \hat{R}_{t'}(f) + \frac{18 \ln(1/\delta_t)}{t'}$$

and so f^* is not eliminated in any elimination step and remains in F_t .

Furthermore, suppose f is also not eliminated and survives in F_t . Then we must have $\hat{R}_t(f) - \hat{R}_t(f^*) \leq 18C^2/t$, or in other words, $Z \leq 18C^2$. Thus, $(X - 4C)^2 \leq 36C^2$, which implies that $X^2 \leq 100C^2$, and hence:

$$\sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}] \leq 100 \ln(1/\delta_t). \quad (4.1)$$

By Lemma 4.3 and since P_t is measurable with respect to the past sigma field up to time $t - 1$, for all $t' \leq t$ we have

$$\mathbf{E}_{x,r} [r(\pi_f(x)) - r(\pi_{f^*}(x))]^2 \leq 2K \mathbf{E}_{x_t', r_t', a_t'} [Y_{t'}].$$

Summing up over all $t' \leq t$, and using (4.1) along with Jensen's inequality we get that

$$\mathbf{E}_{x,r} [r(\pi_f(x)) - r(\pi_{f^*}(x))] \leq \sqrt{\frac{200K \ln(1/\delta_t)}{t}}. \quad \square$$

Lemma 4.2. Fix a function $f \in F$. Suppose we sample x, r from the data distribution D , and an action a from an arbitrary distribution such that r and a are conditionally independent given x . Define the random variable

$$Y = (f(x, a) - r(a))^2 - (f^*(x, a) - r(a))^2.$$

Then we have

$$\begin{aligned} \mathbf{E}_{x,r,a} [Y] &= \mathbf{E}_{x,a} [(f(x, a) - f^*(x, a))^2] \\ \mathbf{Var}_{x,r,a} [Y] &\leq 4 \mathbf{E}_{x,r,a} [Y]. \end{aligned}$$

Proof. Using shorthands f_{xa} for $f(x, a)$ and r_a for $r(a)$, we can rearrange the definition of Y as

$$Y = (f_{xa} - f_{xa}^*)(f_{xa} + f_{xa}^* - 2r_a). \quad (4.2)$$

Hence, we have

$$\begin{aligned} \mathbf{E}_{x,r,a} [Y] &= \mathbf{E}_{x,r,a} [(f_{xa} - f_{xa}^*)(f_{xa} + f_{xa}^* - 2r_a)] \\ &= \mathbf{E}_{x,a} \mathbf{E}_{r|x} [(f_{xa} - f_{xa}^*)(f_{xa} + f_{xa}^* - 2r_a)] \\ &= \mathbf{E}_{x,a} \left[(f_{xa} - f_{xa}^*) \left(f_{xa} + f_{xa}^* - 2 \mathbf{E}_{r|x} [r_a] \right) \right] \\ &= \mathbf{E}_{x,a} [(f_{xa} - f_{xa}^*)^2], \end{aligned}$$

proving the first part of the lemma. From (4.2), noting that f_{xa}, f_{xa}^*, r_a are between 0 and 1, we obtain

$$\begin{aligned} Y^2 &\leq (f_{xa} - f_{xa}^*)^2 (f_{xa} + f_{xa}^* - 2r_a)^2 \\ &\leq 4(f_{xa} - f_{xa}^*)^2, \end{aligned}$$

yielding the second part of the lemma:

$$\begin{aligned} \mathbf{Var}_{x,r,a} [Y] &\leq \mathbf{E}_{x,r,a} [Y^2] \leq 4 \mathbf{E}_{x,r,a} [(f_{xa} - f_{xa}^*)^2] \\ &= 4 \mathbf{E}_{x,r,a} [Y]. \quad \square \end{aligned}$$

Next we show how the random variable Y defined in Lemma 4.2 relates to the regret in a single round:

Lemma 4.3. In the setup of Lemma 4.2, assume further that the action a is sampled from a conditional distribution $p(\cdot|x)$ which satisfies the following constraint, for $f' = f$ and $f' = f^*$:

$$\mathbf{E}_x \left[\frac{1}{p(\pi_{f'}(x)|x)} \right] \leq K. \quad (4.3)$$

Then we have

$$\mathbf{E}_{x,r} [r(\pi_{f^*}(x)) - r(\pi_f(x))]^2 \leq 2K \mathbf{E}_{x,r,a} [Y].$$

This lemma is essentially a refined form of theorem 6.1 in Beygelzimer and Langford (2009) which analyzes the regression approach to learning in contextual bandit settings.

Proof. Throughout, we continue using the shorthand $f_{x,a}$ for $f(x, a)$. Given a context x , let $\tilde{a} = \pi_f(x)$ and $a^* = \pi_{f^*}(x)$. Define the random variable

$$\Delta_x = \mathbf{E}_{r|x} \left[r(\pi_{f^*}(x)) - r(\pi_f(x)) \right] = f_{x a^*}^* - f_{x \tilde{a}}^* .$$

Note that $\Delta_x \geq 0$ because f^* prefers a^* over \tilde{a} for context x . Also we have $f_{x \tilde{a}} \geq f_{x a^*}$ since f prefers \tilde{a} over a^* for context x . Thus,

$$f_{x \tilde{a}} - f_{x \tilde{a}}^* + f_{x a^*}^* - f_{x a^*} \geq \Delta_x . \quad (4.4)$$

As in proof of Lemma 4.2,

$$\begin{aligned} \mathbf{E}_{r,a|x} [Y] &= \mathbf{E}_{a|x} [(f_{x a} - f_{x a}^*)^2] \\ &\geq p(\tilde{a}|x)(f_{x \tilde{a}} - f_{x \tilde{a}}^*)^2 + p(a^*|x)(f_{x a^*}^* - f_{x a^*})^2 \\ &\geq \frac{p(\tilde{a}|x)p(a^*|x)}{p(\tilde{a}|x) + p(a^*|x)} \Delta_x^2 . \end{aligned} \quad (4.5)$$

The last inequality follows by first applying the chain

$$\underline{a}x^2 + \underline{b}y^2 = \frac{ab(\underline{x} + \underline{y})^2 + (\underline{a}x - \underline{b}y)^2}{\underline{a} + \underline{b}} \geq \frac{ab}{\underline{a} + \underline{b}} (\underline{x} + \underline{y})^2$$

(valid for $\underline{a}, \underline{b} > 0$), and then applying inequality (4.4).

For convenience, define

$$Q_x = \frac{p(\tilde{a}|x)p(a^*|x)}{p(\tilde{a}|x) + p(a^*|x)}, \quad \text{i.e.,} \quad \frac{1}{Q_x} = \frac{1}{p(\tilde{a}|x)} + \frac{1}{p(a^*|x)} .$$

Now, since p satisfies the constraint (4.3) for $f' = f$ and $f' = f^*$, we conclude that

$$\mathbf{E}_x \left[\frac{1}{Q_x} \right] = \mathbf{E}_x \left[\frac{1}{p(\tilde{a}|x)} \right] + \mathbf{E}_x \left[\frac{1}{p(a^*|x)} \right] \leq 2K . \quad (4.6)$$

We now have

$$\begin{aligned} \mathbf{E}_x [\Delta_x]^2 &= \mathbf{E}_x \left[\frac{1}{\sqrt{Q_x}} \cdot \sqrt{Q_x} \Delta_x \right]^2 \\ &\leq \mathbf{E}_x \left[\frac{1}{Q_x} \right] \mathbf{E}_x [Q_x \Delta_x^2] \\ &\leq 2K \mathbf{E}_{x,r,a} [Y] , \end{aligned}$$

where the first inequality follows from the Cauchy-Schwarz inequality and the second from the inequalities (4.5) and (4.6). \square

5 Lower bound

Here we prove a lower bound showing that the realizability assumption is *not* enough in general to eliminate a dependence on the number of actions K . The structure of this proof is similar to an earlier lower bound (Auer et al., 2003) differing in two ways: it applies to regressors of the sort we consider, and we work N , the number of regressors, into the lower bound. Since for every policy there exists a regressor with argmax on that regressor realizing the policy, this lower bound also applies to policy based algorithms.

Theorem 5.1. *For every N and K such that $\ln N / \ln K \leq T$, and every algorithm \mathbf{A} , there exists a function class F of cardinality at most N and a distribution $D(x, r)$ for which the realizability assumption holds, but the expected regret of \mathbf{A} is $\Omega(\sqrt{KT \ln N / \ln K})$.*

Proof. Instead of directly selecting F and D for which the expected regret of \mathbf{A} is $\Omega(\sqrt{KT \ln N / \ln K})$, we create a distribution over instances (F, D) and show that the expected regret of \mathbf{A} is $\Omega(\sqrt{KT \ln N / \ln K})$ when the expectation is taken also over our choice of the instance. This will immediately yield a statement of the theorem, since the algorithm must suffer at least this amount of regret on one of the instances.

The proof proceeds via a reduction to the construction used in the lower bound of Theorem 5.1 of Auer et al. (2003). We will use M different contexts for a suitable number M . To define the regressor class F , we begin with the policy class G consisting of all the K^M mappings of the form $g : \mathcal{X} \rightarrow \mathcal{A}$, where $\mathcal{X} = \{1, 2, \dots, M\}$ and $\mathcal{A} = \{1, 2, \dots, K\}$. We require M to be the largest integer such that $K^M \leq N$, i.e., $M = \lfloor \ln N / \ln K \rfloor$. Each mapping $g \in G$ defines a regressor $f_g \in F$ as follows:

$$f_g(x, a) = \begin{cases} 1/2 + \epsilon & \text{if } a = g(x) \\ 1/2 & \text{otherwise.} \end{cases}$$

The rewards are generated by picking a function $f \in F$ uniformly at random at the beginning. Equivalently, we choose a mapping g that independently maps each context $x \in \mathcal{X}$ to a random action $a \in \mathcal{A}$, and set $f = f_g$. In each round t , a context x_t is picked uniformly from \mathcal{X} . For any action a , a reward $r_t(a)$ is generated as a $\{0, 1\}$ Bernoulli trial with probability of 1 being equal to $f(x, a)$.

Now fix a context $x \in \mathcal{X}$. We condition on all of the randomness of the algorithm \mathbf{A} , the choices of the contexts x_t for $t = 1, 2, \dots, T$, and the values of $g(x')$ for $x' \neq x$. Thus the only randomness left is in the choice of $g(x)$ and the realization of the rewards in each round. Let \mathbf{P}' denote the reward distribution where the rewards of any action a for context x are chosen to be $\{0, 1\}$ uniformly at random (the rewards for other contexts $x' \neq x$ are still chosen according to $f(x', a)$, however), and let \mathbf{E}' denote the expectation under \mathbf{P}' .

Let T_x be the rounds t where the context x_t is x . Now fix an action $a \in \mathcal{A}$ and let S_a be a random variable denoting the number of rounds $t \in T_x$ when \mathbf{A} chooses $a_t = a$. Note that conditioned on $g(x) = a$, the random variable S_a counts the number of rounds in T_x that \mathbf{A} chooses the optimal action a .

We use a corollary of Lemma A.1 in Auer et al. (2003):

Corollary 5.1 (Auer et al., 2003). *Conditioned on the choices of the contexts x_t for $t = 1, 2, \dots, T$, and the val-*

ues of $g(x')$ for $x' \neq x$, we have

$$\mathbf{E}[S_a | g(x) = a] \leq \mathbf{E}'[S_a] + |T_x| \sqrt{2\epsilon^2 \mathbf{E}'[S_a]}.$$

The proof uses the fact that when $g(x) = a$, rewards chosen using \mathbf{P}' are identical to those from the true distribution except for the rounds when \mathbf{A} chooses the action a .

Thus, if N_x is a random variable that counts the number the rounds in T_x that \mathbf{A} chooses the optimal action for x (without conditioning on $g(x)$), we have

$$\begin{aligned} \mathbf{E}[N_x] &= \mathbf{E}_{g(x)} [\mathbf{E}[S_{g(x)}]] \\ &\leq \mathbf{E}_{g(x)} \left[\mathbf{E}'[S_{g(x)}] + |T_x| \sqrt{2\epsilon^2 \mathbf{E}'[S_{g(x)}]} \right] \\ &\leq \mathbf{E}_{g(x)} \left[\mathbf{E}'[S_{g(x)}] \right] + |T_x| \sqrt{2\epsilon^2 \mathbf{E}_{g(x)} [\mathbf{E}'[S_{g(x)}]]}, \end{aligned}$$

by Jensen's inequality. Now note that

$$\begin{aligned} \mathbf{E}_{g(x)} [\mathbf{E}'[S_{g(x)}]] &= \mathbf{E}_{g(x)} \left[\mathbf{E}' \left[\sum_{t \in T_x} \mathbf{1}\{a_t = g(x)\} \right] \right] \\ &= \sum_{t \in T_x} \mathbf{E}' [\mathbf{E}_{g(x)} [\mathbf{1}\{a_t = g(x)\}]] \\ &= \sum_{t \in T_x} \mathbf{E}' \left[\frac{1}{K} \right] = \frac{|T_x|}{K}. \end{aligned}$$

The third equality follows because $g(x)$ is independent of the choices of the contexts x_t for $t = 1, 2, \dots, T$, and $g(x')$ for $x' \neq x$, and its distribution is uniform on \mathcal{A} . Thus

$$\mathbf{E}[N_x] \leq \frac{|T_x|}{K} + |T_x| \sqrt{2\epsilon^2 \frac{|T_x|}{K}}.$$

Since in the rounds in $T_x \setminus N_x$, the algorithm \mathbf{A} suffers an expected regret of ϵ , the expected regret of \mathbf{A} over all the rounds in T_x is at least $\Omega\left(\epsilon|T_x| - \frac{\epsilon^2}{\sqrt{K}}|T_x|^{3/2}\right)$. Note that this lower bound is independent of the choice of $g(x')$ for $x' \neq x$. Thus, we can remove the conditioning on $g(x')$ for $x' \neq x$ and conclude that only conditioned on the choices of the contexts x_t for $t = 1, 2, \dots, T$, the expected regret of the algorithm over all the rounds in T_x is at least $\Omega\left(\epsilon|T_x| - \frac{\epsilon^2}{\sqrt{K}}|T_x|^{3/2}\right)$. Summing up over all x , and removing the conditioning on the choices of the contexts x_t for $t = 1, 2, \dots, T$ by taking an expectation, we get the following lower bound on the expected regret of \mathbf{A} :

$$\Omega \left(\sum_{x \in \mathcal{X}} \left(\epsilon \mathbf{E}[|T_x|] - \frac{\epsilon^2}{\sqrt{K}} \mathbf{E}[|T_x|^{3/2}] \right) \right).$$

Note that $|T_x|$ is distributed as Binomial($T, 1/M$). Thus,

$\mathbf{E}[|T_x|] = T/M$. Furthermore, by Jensen's inequality

$$\begin{aligned} \mathbf{E}[|T_x|^{3/2}] &\leq \sqrt{\mathbf{E}[|T_x|^3]} \\ &= \left(\frac{T}{M} + \frac{3T(T-1)}{M^2} + \frac{T(T-1)(T-2)}{M^3} \right)^{1/2} \\ &\leq \frac{\sqrt{5}T^{3/2}}{M^{3/2}}, \end{aligned}$$

as long as $M \leq T$. Plugging these bounds in, the lower bound on the expected regret becomes

$$\Omega \left(\epsilon T - \frac{\epsilon^2}{\sqrt{KM}} T^{3/2} \right).$$

Choosing $\epsilon = \Theta(\sqrt{KM/T})$, we get that the expected regret of \mathbf{A} is lower bounded by

$$\Omega(\sqrt{KMT}) = \Omega(\sqrt{KT \ln N / \ln K}). \quad \square$$

6 Analysis of nontriviality

Since the worst-case regret bound of our new algorithm is the same as for agnostic algorithms, a skeptic could conclude that there is no power in the realizability assumption. Here, we show that in some cases, realizability assumption can be very powerful in reducing regret.

Theorem 6.1. *For any algorithm \mathbf{A} working with a set of policies (rather than regressors), there exists a set of regressors F and a distribution D satisfying the realizability assumption such that the regret of \mathbf{A} using the set Π_F is $\tilde{\Omega}(\sqrt{TK \ln N})$, but the expected regret of Regressor Elimination using F is at most $\mathcal{O}(\ln(N/\delta))$.*

Proof. Let F' be the set of functions and D the data distribution that achieve the lower bound of Theorem 5.1 for the algorithm \mathbf{A} . Using Lemma 6.1 (see below), there exists a set of functions F such that $\Pi_F = \Pi_{F'}$ and the expected regret of Regressor Elimination using F is at most $\mathcal{O}(\ln(N/\delta))$. This set of functions F and distribution D satisfy the requirements of the theorem. \square

Lemma 6.1. *For any distribution D and a set of policies Π containing the optimal policy, there exists a set of functions F satisfying the realizability assumption, such that $\Pi = \Pi_F$ and the regret of regressor elimination using F is at most $\mathcal{O}(\ln(N/\delta))$.*

Proof. The idea is to build a set of functions F such that $\Pi = \Pi_F$, and for the optimal policy π^* the corresponding function f^* exactly gives the expected rewards for each context x and a , but for any other policy π the corresponding function f gives a terrible estimate, allowing regressor elimination to eliminate them quickly.

The construction is as follows. For π^* , we define the function f^* as $f^*(x, a) = \mathbf{E}_{x,r}[r(a)]$. By optimality of π^* ,

$\pi_{f^*} = \pi^*$. For every other policy π we construct an f such that $\pi = \pi_f$ but for which $f(x, a)$ is a very bad estimate of $\mathbf{E}_{x,r}[r(a)]$ for all actions a . Fix x and consider two cases: the first is that $\mathbf{E}_{r|x}[r(\pi(x))] > 0.75$ and the other is that $\mathbf{E}_{r|x}[r(\pi(x))] \leq 0.75$. In the first case, we let $f(x, \pi(x)) = 0.51$. In the second case we let $f(x, \pi(x)) = 1.0$. Now consider each other action a' in turn. If $\mathbf{E}_{r|x}[r(a')] > 0.25$ then we let $f(x, a') = 0$, and if $\mathbf{E}_{r|x}[r(a')] \leq 0.25$ we let $f(x, a') = 0.5$.

The regressor elimination algorithm eliminates regressor with a too-large squared loss regret. Now fix any policy $\pi \neq \pi^*$, and the corresponding f , define, as in the proof of Lemma 4.1, the random variable

$$Y_t = (f(x_t, a_t) - r_t(a_t))^2 - (f^*(x_t, a_t) - r_t(a_t))^2.$$

Note that

$$\mathbf{E}_t[Y_t] = \mathbf{E}_{x_t, a_t} [(f(x_t, a_t) - f^*(x_t, a_t))^2] \geq \frac{1}{20}, \quad (6.1)$$

since for all (x, a) , $(f(x, a) - f^*(x, a))^2 \geq \frac{1}{20}$ by construction. This shows that the expected regret is significant.

Now suppose f is not eliminated and remains in F_t . Then by equation 4.1 we get:

$$\frac{t}{20} \leq \sum_{t'=1}^t \mathbf{E}_{t'}[Y_{t'}] \leq 100 \ln(1/\delta_t).$$

The above bound holds with probability $1 - \delta_t N t \log_2(t)$ uniformly for all $f \in F_t$. Using the choice of $\delta_t = \delta / 2Nt^3 \log_2(t)$, we note that the bound fails to hold when $t > 10^6 \ln(N/\delta)$. Thus, within $10^6 \ln(N/\delta)$ rounds all suboptimal regressors are eliminated, and the algorithm suffers no regret thereafter. Since the rewards are bounded in $[0, 1]$, the total regret in the first $10^6 \ln(N/\delta)$ rounds can be at most $10^6 \ln(N/\delta)$, giving us the desired bound. \square

7 Removing the dependence on D

While Algorithm 1 is conceptually simple and enjoys nice theoretical guarantees, it has a serious drawback that it depends on the distribution D from which the contexts x_t 's are drawn in order to specify the constraint (3.1). A similar issue was faced in the earlier work of Dudík et al. (2011), where they replace the expectation under D with a sample average over the contexts observed. We now discuss a similar modification for Algorithm 1 and give a sketch of the regret analysis.

The key change in Algorithm 1 is to replace the constraint (3.1) with the sample version. Let $H_t = \{x_1, x_2, \dots, x_{t-1}\}$, and denote by $x \sim H_t$ the act of selecting a context x from H_t uniformly at random. Now we

pick a distribution P_t on F_{t-1} such that

$$\forall f \in F_{t-1} : \mathbf{E}_{x \sim H_t} \left[\frac{1}{P'_t(\pi_f(x)|x)} \right] \leq \mathbf{E}_{x \sim H_t} [|A(F_{t-1}, x)|] \quad (7.1)$$

Since Lemma A.1 applies to any distribution on the contexts, in particular, the uniform distribution on H_t , this constraint is still feasible. To justify this sample based approximation, we appeal to Theorem 6 of Dudík et al. (2011) which shows that for any $\epsilon \in (0, 1)$ and $t \geq 16K \ln(8KN/\delta)$, with probability at least $1 - \delta$

$$\begin{aligned} & \mathbf{E}_{x \sim D} \left[\frac{1}{P'_t(\pi_f(x)|x)} \right] \\ & \leq (1 + \epsilon) \mathbf{E}_{x \sim H_t} \left[\frac{1}{P'_t(\pi_f(x)|x)} \right] + \frac{7500}{\epsilon^3} K. \end{aligned}$$

Using Equation (7.1), since $|A(F_{t-1}, x_t)| \leq K$, we get

$$\mathbf{E}_{x \sim D} \left[\frac{1}{P'_t(\pi_f(x)|x)} \right] \leq 7525K,$$

using $\epsilon = 0.999$. The remaining analysis of the algorithm remains the same as before, except we now apply Lemma 4.3 with a worse constant in the condition (4.3).

8 Conclusion

The included results gives us a basic understanding of the realizable assumption setting: it can, but does not necessarily, improve our ability to learn.

We did not address computational complexity in this paper. There are some reasons to be hopeful however. Due to the structure of the realizability assumption, an eliminated regressor continues to have an increasingly poor regret over time, implying that it may be possible to avoid the elimination step and simply restrict the set of regressors we care about when constructing a distribution. A basic question then is: can we make the formation of this distribution computationally tractable?

Another question for future research is the extension to infinite function classes. One would expect that this just involves replacing the log cardinality with something like a metric entropy or Rademacher complexity of F . This is not completely immediate since we are dealing with martingales, and direct application of covering arguments seems to yield a suboptimal $\mathcal{O}(1/\sqrt{t})$ rate in Lemma 4.1. Extending the variance based bound coming from Freedman's inequality from a single martingale to a supremum over function classes would need a Talagrand-style concentration inequality for martingales which is not available in the literature to the best of our knowledge. Understanding this issue better is an interesting topic for future work.

Acknowledgements This research was done while AA, SK and RES were visiting Yahoo!.

References

- P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77, 2003.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *J. Mach. Learn. Res.*, 3:397–422, March 2003.
- P. L. Bartlett, V. Dani, T. P. Hayes, S. Kakade, A. Rakhlin, and A. Tewari. High-probability regret bounds for bandit online linear optimization. In *COLT*, 2008.
- A. Beygelzimer and J. Langford. The offset tree for learning with partial labels. In *KDD*, 2009.
- A. Beygelzimer, J. Langford, L. Li, L. Reyzin, and R.E. Schapire. An optimal high probability algorithm for the contextual bandit problem. *CoRR*, 2010. URL <http://arxiv.org/abs/1002.4058>.
- W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In *AISTATS*, 2011.
- V. Dani, T. P. Hayes, and S. M. Kakade. Stochastic linear optimization under bandit feedback. In *Proceedings of the 21st Annual Conference on Learning Theory*, 2008.
- M. Dudík, D. Hsu, S. Kale, N. Karampatziakis, J. Langford, L. Reyzin, and T. Zhang. Efficient optimal learning for contextual bandits. In *UAI*, 2011.
- S. Filippi, O. Cappé, A. Garivier, and Cs. Szepesvári. Parametric bandits: The generalized linear case. In *NIPS*, 2010.
- T. L. Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
- Maurice Sion. On general minimax theorems. *Pacific J. Math.*, 8(1):171–176, 1958.

A Feasibility

Lemma A.1. *There exists a distribution P_t on F_{t-1} satisfying the constraint (3.1).*

Proof. Let Δ_{t-1} refer to the space of all distributions on F_{t-1} . We observe that Δ_{t-1} is a convex, compact set. For a distribution $Q \in \Delta_{t-1}$, define the conditional distribution $\tilde{Q}(\cdot|x)$ on \mathcal{A} as sample $f \sim Q$, and return $\pi_f(x)$. Note that $Q'(a|x) = (1 - \mu)\tilde{Q}(a|x) + \mu/K_x$, where $K_x := |A(F_{t-1}, x)|$ for notational convenience.

The feasibility of constraint (3.1) can be written as

$$\min_{P_t \in \Delta_{t-1}} \max_{f \in F_{t-1}} \mathbf{E}_x \left[\frac{1}{P'_t(\pi_f(x)|x)} \right] \leq \mathbf{E}_x [|A(F_{t-1}, x)|].$$

The LHS is equal to

$$\min_{P_t \in \Delta_{t-1}} \max_{Q \in \Delta_{t-1}} \mathbf{E}_x \left[\sum_{f \in F_{t-1}} \frac{Q(f)}{P'_t(\pi_f(x)|x)} \right],$$

where we recall that P'_t is the distribution induced on \mathcal{A} by P_t as before. The function

$$\mathbf{E}_x \left[\sum_{f \in F_{t-1}} \frac{Q(f)}{P'_t(\pi_f(x)|x)} \right]$$

is linear (and hence concave) in Q and convex in P_t . Applying Sion’s Minimax Theorem (stated below as Theorem A.1), we see that the LHS is equal to

$$\begin{aligned} & \max_{Q \in \Delta_{t-1}} \min_{P_t \in \Delta_{t-1}} \mathbf{E}_x \left[\sum_{f \in F_{t-1}} \frac{Q(f)}{P'_t(\pi_f(x)|x)} \right] \\ & \leq \max_{Q \in \Delta_{t-1}} \mathbf{E}_x \left[\sum_{f \in F_{t-1}} \frac{Q(f)}{Q'(\pi_f(x)|x)} \right] \\ & = \max_{Q \in \Delta_{t-1}} \mathbf{E}_x \left[\sum_{a \in A(F_{t-1}, x)} \sum_{f \in F_{t-1} : \pi_f(x) = a} \frac{Q(f)}{Q'(a|x)} \right] \\ & = \max_{Q \in \Delta_{t-1}} \mathbf{E}_x \left[\sum_{a \in A(F_{t-1}, x)} \frac{\tilde{Q}(a|x)}{Q'(a|x)} \right] \\ & = \max_{Q \in \Delta_{t-1}} \mathbf{E}_x \left[\frac{1}{1 - \mu} \cdot \sum_{a \in A(F_{t-1}, x)} \left[1 - \frac{\mu}{K_x Q'(a|x)} \right] \right] \\ & \leq \max_{Q \in \Delta_{t-1}} \mathbf{E}_x [K_x] . \end{aligned}$$

The last inequality uses the fact that for any distribution P on $\{1, 2, \dots, K\}$, $\sum_{i=1}^K [1/P(i)]$ is minimized when all $P(i)$ equal $1/K$. Hence the constraint is always feasible. \square

Theorem A.1 (see Theorem 3.4 of Sion, 1958). *Let U and V be compact and convex sets, and $\phi : U \times V \rightarrow \mathbb{R}$ a function which for all $v \in V$ is convex and continuous in u and for all $u \in U$ is concave and continuous in v . Then*

$$\min_{u \in U} \max_{v \in V} \phi(u, v) = \max_{v \in V} \min_{u \in U} \phi(u, v) .$$

B Freedman-style Inequality

Lemma B.1 (see Bartlett et al., 2008). *Suppose X_1, X_2, \dots, X_T is a martingale difference sequence with $|X_t| \leq b$ for all t . Let $V = \sum_{t=1}^T \mathbf{Var}_t[X_t]$ be the sum of conditional variances. Then for any $\delta < 1/e^2$, with probability at least $1 - \log_2(T)\delta$ we have*

$$\sum_{t=1}^T X_t \leq 4\sqrt{V \ln(1/\delta)} + 2b \ln(1/\delta).$$