# PAC-Bayesian Analysis of Contextual Bandits 

Yevgeny Seldin ${ }^{1,4}$ Peter Auer ${ }^{2}$ François Laviolette ${ }^{3}$ John Shawe-Taylor ${ }^{4}$ Ronald Ortner ${ }^{2}$<br>${ }^{1}$ Max Planck Institute for Intelligent Systems, Tübingen, Germany<br>${ }^{2}$ Chair for Information Technology, Montanuniversität Leoben, Austria<br>${ }^{3}$ Département d'informatique, Université Laval, Québec, Canada<br>${ }^{4}$ Department of Computer Science, University College London, UK<br>seldin@tuebingen.mpg.de, \{auer, ronald.ortner\}@unileoben.ac.at, francois.laviolette@ift.ulaval.ca, jst@cs.ucl.ac.uk


#### Abstract

We derive an instantaneous (per-round) data-dependent regret bound for stochastic multiarmed bandits with side information (also known as contextual bandits). The scaling of our regret bound with the number of states (contexts) $N$ goes as $\sqrt{N I_{\rho_{t}}(S ; A)}$, where $I_{\rho_{t}}(S ; A)$ is the mutual information between states and actions (the side information) used by the algorithm at round $t$. If the algorithm uses all the side information, the regret bound scales as $\sqrt{N \ln K}$, where $K$ is the number of actions (arms). However, if the side information $I_{\rho_{t}}(S ; A)$ is not fully used, the regret bound is significantly tighter. In the extreme case, when $I_{\rho_{t}}(S ; A)=0$, the dependence on the number of states reduces from linear to logarithmic. Our analysis allows to provide the algorithm large amount of side information, let the algorithm to decide which side information is relevant for the task, and penalize the algorithm only for the side information that it is using de facto. We also present an algorithm for multiarmed bandits with side information with $O(K)$ computational complexity per game round.


## 1 Introduction

Multiarmed bandits with side information are an elegant mathematical model for many real-life interactive systems, such as personalized online advertising, personalized medical treatment, and so on. This model is also known as contextual bandits or associative bandits (Kaelbling, 1994, Strehl et al., 2006, Langford and Zhang, 2007, Beygelzimer et al., 2011). In multiarmed bandits with side information the learner repeatedly observes states (side information) $\left\{s_{1}, s_{2}, \ldots\right\}$ (for example, symptoms of a patient) and has to perform actions (for example, prescribe drugs), such that the expected regret is minimized. The regret is usually measured by the difference between the reward that could be achieved by the best (unknown) fixed policy (for example, the number of patients that would be cured if we knew the best drug for each set of symptoms) and the reward obtained by the algorithm (the number of patients that were actually cured).

Most of the existing analyses of multiarmed bandits with side information has focused on the adversarial (worst-case) model, where the sequence of rewards associated with each state-action pair is chosen by an adversary. However, many problems in real-life are not adversarial. We derive datadependent analysis for stochastic multiarmed bandits with side information. In the stochastic setting the rewards for each state-action pair are drawn from a fixed unknown distribution. The sequence of states is also drawn from a fixed unknown distribution. We restrict ourselves to problems with finite number of states $N$ and finite number of actions $K$ and leave generalization to continuous state and action spaces to future work. We also do not assume any structure of the state space. Thus, for us a state is just a number between 1 and $N$. For example, in online advertising the state can be the country from which a web page is accessed.

The result presented in this paper exhibits adaptive dependency on the side information (state identity) that is actually used by the algorithm. This allows us to provide the algorithm a large amount of side information and let the algorithm decide, which of this side information is actually relevant to the task. For example, in online advertising we can increase the state resolution and provide the algorithm the town from which the web page was accessed, but if this refined state information is not used by the algorithm the regret bound will not deteriorate. This can be opposed to existing analysis of adversarial multiarmed bandits, where the regret bound depends on a predefined complexity of the underlying expert class (Beygelzimer et al., 2011). Thus, the existing analysis of adversarial multiarmed bandits would either become looser if we add more side information or a-priori limit the usage of the side information through its internal structure. (We note that through the relation between PAC-Bayesian analysis and the analysis of adversarial online learning described in Banerjee (2006) it might be possible to extend our analysis to adversarial setting, but we leave this research direction to future work.)

The idea of regularization by relevant mutual information goes back to the Information Bottleneck principle in supervised and unsupervised learning (Tishby et al., 1999). Tishby and Polani (2010) further suggested to measure the complexity of a policy in reinforcement learning by the mutual information between states and actions used by the policy. We note, however, that our starting point is the regret bound and we derive the regularization term from our analysis without introducing it a-priori. The analysis also provides time and data dependent weighting of the regularization term.

Our results are based on PAC-Bayesian analysis (Shawe-Taylor and Williamson, 1997, ShaweTaylor et al., 1998, McAllester, 1998, Seeger, 2002), which was developed for supervised learning within the PAC (Probably Approximately Correct) learning framework (Valiant, 1984). In PACBayesian analysis the complexity of a model is defined by a user-selected prior over a hypothesis space. Unlike in VC-dimension-based approaches and their successors, where the complexity is defined for a hypothesis class, in PAC-Bayesian analysis the complexity is defined for individual hypotheses. The analysis provides an explicit trade-off between individual model complexity and its empirical performance and a high probability guarantee on the expected performance.
An important distinction between supervised learning and problems with limited feedback, such as multiarmed bandits and reinforcement learning more generally, is the fact that in supervised learning the training set is given, whereas in reinforcement learning the training set is generated by the learner as it plays the game. In supervised learning every hypothesis in a hypothesis class can be evaluated on all the samples, whereas in reinforcement learning rewards of one action cannot be used to evaluate another action. Recently, Seldin et al. (2011b,a) generalized PAC-Bayesian analysis to martingales and suggested a way to apply it under limited feedback. Here, we apply this generalization to multiarmed bandits with side information.
The remainder of the paper is organized as follows. We start with definitions in Section 2 and provide our main results in Section 3, which include an instantaneous regret bound and a new algorithm for stochastic multiarmed bandits with side information. In Section 4 we present an experiment that illustrates our theoretical results. Then, we dive into the proof of our main results in Section 5 and discuss the paper in Section 6.

## 2 Definitions

In this section we provide all essential definitions for our main results in the following section. We start with the definition of stochastic multiarmed bandits with side information. Let $\mathcal{S}$ be a set of $|\mathcal{S}|=N$ states and let $\mathcal{A}$ be a set of $|\mathcal{A}|=K$ actions, such that any action can be performed in any state. Let $s \in \mathcal{S}$ denote the states and $a \in \mathcal{A}$ denote the actions. Let $R(a, s)$ be the expected reward for performing action $a$ in state $s$. At each round $t$ of the game the learner is presented a state $S_{t}$ drawn i.i.d. according to an unknown distribution $p(s)$. The learner draws an action $A_{t}$ according to his choice of a distribution (policy) $\pi_{t}(a \mid s)$ and obtains a stochastic reward $R_{t}$ with expected value $R\left(A_{t}, S_{t}\right)$. Let $\left\{S_{1}, S_{2}, \ldots\right\}$ denote the sequence of observed states, $\left\{\pi_{1}, \pi_{2}, \ldots\right\}$ the sequence of policies played, $\left\{A_{1}, A_{2}, \ldots\right\}$ the sequence of actions played, and $\left\{R_{1}, R_{2}, \ldots\right\}$ the sequence of observed rewards. Let $\mathcal{T}_{t}=\left\{\left\{S_{1}, \ldots, S_{t}\right\},\left\{\pi_{1}, \ldots, \pi_{t}\right\},\left\{A_{1}, \ldots, A_{t}\right\},\left\{R_{1}, \ldots, R_{t}\right\}\right\}$ denote the history of the game up to time $t$.

Assume that $\pi_{t}(a \mid s)>0$ for all $t, a$, and $s$. For $t \geq 1, a \in\{1, \ldots, K\}$, and the sequence of observed states $\left\{S_{1}, \ldots, S_{t}\right\}$ define a set of random variables $R_{t}^{a, S_{t}}$ :

$$
R_{t}^{a, S_{t}}=\left\{\begin{array}{cl}
\frac{1}{\pi_{t}\left(a \mid S_{t}\right)} R_{t}, & \text { if } A_{t}=a \\
0, & \text { otherwise }
\end{array}\right.
$$

(The variables $R_{t}^{a, s}$ are defined only for the observed state $s=S_{t}$.) Note that whenever defined, $\mathbb{E}\left[R_{t}^{a, S_{t}} \mid \mathcal{T}_{t-1}, S_{t}\right]=R\left(a, S_{t}\right)$. The definition of $R_{t}^{a, s}$ is generally known as importance weighted sampling (Sutton and Barto, 1998). Importance weighted sampling is required for application of PAC-Bayesian analysis, as will be shown in the technical part of the paper.

Define $n_{t}(s)=\sum_{\tau=1}^{t} \mathbb{I}_{\left\{S_{\tau}=s\right\}}$ as the number of times state $s$ appeared up to time $t$ ( $\mathbb{I}$ is the indicator function). We define the empirical rewards of state-action pairs as:

$$
\hat{R}_{t}(a, s)=\left\{\begin{array}{cl}
\frac{\sum_{\left\{\tau=1, \ldots, t: S_{\tau}=s\right\}} R_{\tau}^{a, s}}{n_{t}(s)}, & \text { if } n_{t}(s)>0 \\
0, & \text { otherwise }
\end{array}\right.
$$

Note that whenever $n_{t}(s)>0$ we have $\mathbb{E} \hat{R}_{t}(a, s)=R(a, s)$. For every state $s$ we define the "best" action in that state as $a^{*}=\arg \max _{a} R(a, s)$ (if there are multiple "best" actions, one of them is chosen arbitrarily). We then define the expected and empirical regret for performing any other action $a$ in state $s$ as:

$$
\Delta(a, s)=R\left(a^{*}(s), s\right)-R(a, s), \quad \hat{\Delta}_{t}(a, s)=\hat{R}_{t}\left(a^{*}(s), s\right)-\hat{R}_{t}(a, s)
$$

Let $\hat{p}_{t}(s)=\frac{n_{t}(s)}{t}$ be the empirical distribution over states observed up to time $t$. For any policy $\rho(a \mid s)$ we define the empirical reward, empirical regret, and expected regret of the policy as: $\hat{R}_{t}(\rho)=\sum_{s} \hat{p}_{t}(s) \sum_{a} \rho(a \mid s) \hat{R}_{t}(a, s), \hat{\Delta}_{t}(\rho)=\sum_{s} \hat{p}_{t}(s) \sum_{a} \rho(a \mid s) \hat{\Delta}_{t}(a, s)$, and $\Delta(\rho)=$ $\sum_{s} p(s) \sum_{a} \rho(a \mid s) \Delta(a, s)$.
We define the marginal distribution over actions that corresponds to a policy $\rho(a \mid s)$ and the uniform distribution over $\mathcal{S}$ as $\bar{\rho}(a)=\frac{1}{N} \sum_{s} \rho(a \mid s)$ and the mutual information between actions and states corresponding to the policy $\rho(a \mid s)$ and the uniform distribution over $\mathcal{S}$ as

$$
I_{\rho}(S ; A)=\frac{1}{N} \sum_{s, a} \rho(a \mid s) \ln \frac{\rho(a \mid s)}{\bar{\rho}(a)}
$$

For the proof of our main result and also in order to explain the experiments we also have to define a hypothesis space for our problem. This definition is not used in the statement of the main result. Let $\mathcal{H}$ be a hypothesis space, such that each member $h \in \mathcal{H}$ is a deterministic mapping from $\mathcal{S}$ to $\mathcal{A}$. Denote by $a=h(s)$ the action assigned by hypothesis $h$ to state $s$. It is easy to see that the size of the hypothesis space $|\mathcal{H}|=K^{N}$. Denote by $R(h)=\sum_{s \in \mathcal{S}} p(s) R(h(s), s)$ the expected reward of a hypothesis $h$. Define:

$$
\hat{R}_{t}(h)=\frac{1}{t} \sum_{\tau=1}^{t} R_{\tau}^{h\left(S_{\tau}\right), S_{\tau}}
$$

Note that $\mathbb{E} \hat{R}_{t}(h)=R(h)$.
Let $h^{*}=\arg \max _{h \in \mathcal{H}} R(h)$ be the "best" hypothesis (the one that chooses the "best" action in each state). (If there are multiple hypotheses achieving maximal reward pick any of them.) Define:

$$
\Delta(h)=R\left(h^{*}\right)-R(h), \quad \hat{\Delta}_{t}(h)=\hat{R}_{t}\left(h^{*}\right)-\hat{R}_{t}(h)
$$

Any policy $\rho(a \mid s)$ defines a distribution over $\mathcal{H}$ : we can draw an action $a$ for each state $s$ according to $\rho(a \mid s)$ and thus obtain a hypothesis $h \in \mathcal{H}$. We use $\rho(h)$ to denote the respective probability of drawing $h$. For a policy $\rho$ we define $\Delta(\rho)=\mathbb{E}_{\rho(h)}[\Delta(h)]$ and $\hat{\Delta}_{t}(\rho)=\mathbb{E}_{\rho(h)}\left[\hat{\Delta}_{t}(h)\right]$. By marginalization these definitions are consistent with our preceding definitions of $\Delta(\rho)$ and $\hat{\Delta}_{t}(\rho)$.
Finally, let $n_{h}(a)=\sum_{s=1}^{N} \mathbb{I}_{h(s)=a}$ be the number of states in which action $a$ is played by the hypothesis $h$. Let $A^{h}=\left\{\frac{n_{h}(a)}{N}\right\}_{a \in \mathcal{A}}$ be the normalized cardinality profile (histogram) over the
actions played by hypothesis $h$ (with respect to the uniform distribution over $\mathcal{S}$ ). Let $H\left(A^{h}\right)=$ $-\sum_{a} \frac{n_{h}(a)}{N} \ln \frac{n_{h}(a)}{N}$ be the entropy of this cardinality profile. In other words, $H\left(A^{h}\right)$ is the entropy of an action choice of hypothesis $h$ (with respect to the uniform distribution over $\mathcal{S}$ ). Note, that the optimal policy $\rho^{*}(a \mid s)$ (the one, that selects the "best" action in each state) is deterministic and we have $I_{\rho^{*}}(S ; A)=H\left(A^{h^{*}}\right)$.

## 3 Main Results

Our main result is a data and complexity dependent regret bound for a general class of prediction strategies of a smoothed exponential form. Let $\rho_{t}(a)$ be an arbitrary distribution over actions, let

$$
\begin{equation*}
\rho_{t}^{e x p}(a \mid s)=\frac{\rho_{t}(a) e^{\gamma_{t} \hat{R}_{t}(a, s)}}{Z\left(\rho_{t}^{e x p}, s\right)} \tag{1}
\end{equation*}
$$

where $Z\left(\rho_{t}^{\exp }, s\right)=\sum_{a} \rho_{t}(a) e^{\gamma_{t} \hat{R}_{t}(a, s)}$ is a normalization factor, and let

$$
\begin{equation*}
\tilde{\rho}_{t}^{e x p}(a \mid s)=\left(1-K \varepsilon_{t+1}\right) \rho_{t}^{e x p}(a \mid s)+\varepsilon_{t+1} \tag{2}
\end{equation*}
$$

be a smoothed exponential policy. The following theorem provides a regret bound for playing $\tilde{\rho}_{t}^{\text {exp }}$ at round $t+1$ of the game. For generality, we assume that rounds $1, \ldots, t$ were played according to arbitrary policies $\pi_{1}, \ldots, \pi_{t}$.
Theorem 1. Assume that in game rounds $1, \ldots, t$ policies $\left\{\pi_{1}, \ldots, \pi_{t}\right\}$ were played and assume that $\min _{a, s} \pi_{t}(a \mid s) \geq \varepsilon_{t}$ for an arbitrary $\varepsilon_{t}$ that is independent of $\mathcal{T}_{t}$. Let $\rho_{t}(a)$ be an arbitrary distribution over $\mathcal{A}$ that can depend on $\mathcal{T}_{t}$ and satisfies $\min _{a} \rho_{t}(a) \geq \epsilon_{t}$. Let $c>1$ be an arbitrary number that is independent of $\mathcal{T}_{t}$. Then, with probability greater than $1-\delta$ over $\mathcal{T}_{t}$, simultaneously for all policies $\tilde{\rho}_{t}^{\text {exp }}$ defined by (2) that satisfy

$$
\begin{equation*}
\frac{N I_{\rho_{t}^{e x p}}(S ; A)+K(\ln N+\ln K)+\ln \frac{2 m_{t}}{\delta}}{2(e-2) t} \leq \frac{\varepsilon_{t}}{c^{2}} \tag{3}
\end{equation*}
$$

we have:

$$
\begin{equation*}
\Delta\left(\tilde{\rho}_{t}^{\text {exp }}\right) \leq(1+c) \sqrt{\frac{2(e-2)\left(N I_{\rho_{t}^{e x p}}(S ; A)+K(\ln N+\ln K)+\ln \frac{2 m_{t}}{\delta}\right)}{t \varepsilon_{t}}}+\frac{\ln \frac{1}{\epsilon_{t+1}}}{\gamma_{t}}+K \varepsilon_{t+1} \tag{4}
\end{equation*}
$$

where $m_{t}=\ln \left(\sqrt{\frac{(e-2) t}{\ln \frac{2}{\delta}}}\right) / \ln (c)$, and for all $\rho_{t}^{e x p}$ that do not satisfy (3), with the same probability:

$$
\Delta\left(\tilde{\rho}_{t}^{e x p}\right) \leq \frac{2\left(N I_{\rho_{t}^{e x p}}(S ; A)+K(\ln N+\ln K)+\ln \frac{2 m_{t}}{\delta}\right)}{t \varepsilon_{t}}+\frac{\ln \frac{1}{\epsilon_{t+1}}}{\gamma_{t}}+K \varepsilon_{t+1}
$$

Note that the mutual information in Theorem 1 is calculated with respect to $\rho_{t}^{e x p}$ and not $\tilde{\rho}_{t}^{\text {exp }}$. Theorem 1 allows to tune the learning rate $\gamma_{t}$ based on the sample. It also provides an instantaneous regret bound for any algorithm that plays the policies $\left\{\tilde{\rho}_{1}^{\text {exp }}, \tilde{\rho}_{2}^{\text {exp }}, \ldots\right\}$ throughout the game. In order to obtain such a bound we just have to take a decreasing sequence $\left\{\varepsilon_{1}, \varepsilon_{2}, \ldots\right\}$ and substitute $\delta$ in Theorem 1 with $\delta_{t}=\frac{\delta}{t(t+1)}$. Then, by the union bound, the result holds with probability greater than $1-\delta$ for all rounds of the game simultaneously. This leads to Algorithm 1 for stochastic multiarmed bandits with side information. Note that each round of the algorithm takes $O(K)$ time.

Theorem 1 is based on the following regret decomposition and the subsequent theorem and two lemmas that bound the three terms in the decomposition.

$$
\begin{equation*}
\Delta\left(\tilde{\rho}_{t}^{e x p}\right)=\left[\Delta\left(\rho_{t}^{e x p}\right)-\hat{\Delta}_{t}\left(\rho_{t}^{e x p}\right)\right]+\hat{\Delta}_{t}\left(\rho_{t}^{e x p}\right)+\left[R\left(\rho_{t}^{e x p}\right)-R\left(\tilde{\rho}_{t}^{e x p}\right)\right] \tag{5}
\end{equation*}
$$

Theorem 2. Under the conditions of Theorem 1 on $\left\{\pi_{1}, \ldots, \pi_{t}\right\}$ and $c$, simultaneously for all policies $\rho$ that satisfy (3) with probability greater than $1-\delta$ :

$$
\begin{equation*}
\left|\Delta(\rho)-\hat{\Delta}_{t}(\rho)\right| \leq(1+c) \sqrt{\frac{2(e-2)\left(N I_{\rho(S ; A)}+K(\ln N+\ln K)+\ln \frac{2 m_{t}}{\delta}\right)}{t \varepsilon_{t}}} \tag{6}
\end{equation*}
$$

```
Algorithm 1: Algorithm for stochastic contextual bandits. (See text for definitions of \(\varepsilon_{t}\) and \(\gamma_{t}\).)
Input: N, K
\(\hat{R}(a, s) \leftarrow 0\) for all \(a, s\) (These are cumulative [unnormalized] rewards)
\(\rho(a) \leftarrow \frac{1}{K}\) for all \(a\)
\(n(s) \leftarrow 0\) for all \(s\)
\(t \leftarrow 1\)
while not terminated do
    Observe state \(S_{t}\).
    if \(\left(\varepsilon_{t} \geq \frac{1}{K}\right)\) or \(\left(n\left(S_{t}\right)=0\right)\) then
        \(\rho\left(a \mid S_{t}\right) \leftarrow \rho(a)\) for all \(a\)
    else
        \(\rho\left(a \mid S_{t}\right) \leftarrow\left(1-K \varepsilon_{t}\right) \frac{\rho(a) e^{\gamma_{t} \hat{R}\left(a, S_{t}\right) / n\left(S_{t}\right)}}{\sum_{a^{\prime}} \rho\left(a^{\prime}\right) e^{\gamma} \hat{R}\left(a^{\prime}, S_{t}\right) / n\left(S_{t}\right)}+\varepsilon_{t}\) for all \(a\)
        \(\rho(a) \leftarrow \frac{N-1}{N} \rho(a)+\frac{1}{N} \rho\left(a \mid S_{t}\right)\) for all \(a\)
```

    Draw action \(A_{t}\) according to \(\rho\left(a \mid S_{t}\right)\) and play it.
    Observe reward \(R_{t}\).
    \(n\left(S_{t}\right) \leftarrow n\left(S_{t}\right)+1\)
    \(\hat{R}\left(A_{t}, S_{t}\right) \leftarrow \hat{R}\left(A_{t}, S_{t}\right)+\frac{R_{t}}{\rho\left(A_{t} \mid S_{t}\right)}\)
    \(t \leftarrow t+1\)
    and for all $\rho$ that do not satisfy (3) with the same probability:

$$
\left|\Delta(\rho)-\hat{\Delta}_{t}(\rho)\right| \leq \frac{2\left(N I_{\rho}(S ; A)+K(\ln N+\ln K)+\ln \frac{2 m_{t}}{\delta}\right)}{t \varepsilon_{t}}
$$

Note that Theorem 2 holds for all possible $\rho$-s, including those that do not have an exponential form.
Lemma 1. For any distribution $\rho_{t}^{\text {exp }}$ of the form (1), where $\rho_{t}(a) \geq \epsilon$ for all $a$, we have:

$$
\hat{\Delta}_{t}\left(\rho_{t}^{e x p}\right) \leq \frac{\ln \frac{1}{\epsilon}}{\gamma_{t}}
$$

Lemma 2. Let $\tilde{\rho}$ be an $\varepsilon$-smoothed version of a policy $\rho$, such that $\tilde{\rho}(a \mid s)=(1-K \varepsilon) \rho(a \mid s)+\varepsilon$, then

$$
R(\rho)-R(\tilde{\rho}) \leq K \varepsilon
$$

Proof of Theorem 2 is provided in Section 5 and proofs of Lemmas 1 and 2 are provided in the supplementary material.
Comments on Theorem 1. Theorem 1 exhibits what we were looking for: the regret of a policy $\tilde{\rho}_{t}^{e x p}$ depends on the trade-off between its complexity, $N I_{\rho_{t}^{e x p}}(S ; A)$, and the empirical regret, which is bounded by $\frac{1}{\gamma_{t}} \ln \frac{1}{\epsilon_{t+1}}$. We note that $0 \leq I_{\rho_{t}}(S ; A) \leq \ln K$, hence, the result is interesting when $N \gg K$, since otherwise $K \ln K$ term in the bound neutralizes the advantage we get from having small mutual information values. The assumption that $N \gg K$ is reasonable for many applications.

We believe that the dependence of the first term of the regret bound (4) on $\varepsilon_{t}$ is an artifact of our crude upper bound on the variance of the sampling process (given in Lemma 3 in the proof of Theorem 2) and that this term should not be in the bound. This is supported by an empirical study of stochastic multiarmed bandits (Seldin et al., 2011a). With the current bound the best choice for $\varepsilon_{t}$ is $\varepsilon_{t}=(K t)^{-1 / 3}$, which, by integration over the game rounds, yields $O\left(K^{1 / 3} t^{2 / 3}\right)$ dependence of the cumulative regret on the number of arms and game rounds. However, if we manage to derive a tighter analysis and remove $\varepsilon_{t}$ from the first term in (4), the best choice of $\varepsilon_{t}$ will be $\varepsilon_{t}=(K t)^{-1 / 2}$ and the dependence of the cumulative regret on the number of arms and time horizon will improve to $O\left((K t)^{1 / 2}\right)$. One way to achieve this is to apply EXP3.P-style updates (Auer et al., 2002b), however, Seldin et al. (2011a) empirically show that in stochastic environments EXP3 algorithm of Auer et al. (2002b), which is closely related to Algorithm 1, has significantly better performance. Thus, it is desirable to derive a better analysis for EXP3 algorithm in stochastic environments. We note
that although UCB algorithm for stochastic multiarmed bandits (Auer et al., 2002a) is asymptotically better than the EXP3 algorithm, it is not compatible with PAC-Bayesian analysis and we are not aware of a way to derive a UCB-type algorithm and analysis for multiarmed bandits with side information, whose dependence on the number of states would be better than $O(N \ln K)$. Seldin et al. (2011a) also demonstrate that empirically it takes a large number of rounds until the asymptotic advantage of UCB over EXP3 translates into a real advantage in practice.

It is not trivial to minimize (4) with respect to $\gamma_{t}$ analytically. Generally, higher values of $\gamma_{t}$ decrease the second term of the bound, but also lead to more concentrated policies (conditional distributions) $\rho_{t}^{\text {exp }}(a \mid s)$ and thus higher mutual information values $I_{\rho_{t}^{e x p}}(S ; A)$. A simple way to address this trade-off is to set $\gamma_{t}$ such that the contribution of the second term is as close to the contribution of the first term as possible. This can be approximated by taking the value of mutual information from the previous round (or approximation of the value of mutual information from the previous round). More details on parameter setting for the algorithm are provided in the supplementary material.

Comments on Algorithm 1. By regret decomposition (5) and Theorem 2, regret at round $t+1$ is minimized by a policy $\rho_{t}(a \mid s)$ that minimizes a certain trade-off between the mutual information $I_{\rho}(S ; A)$ and the empirical regret $\hat{R}_{t}(\rho)$. This trade-off is analogical to rate-distortion trade-off in information theory (Cover and Thomas, 1991). Minimization of rate-distortion trade-off is achieved by iterative updates of the following form, which are known as Blahut-Arimoto (BA) algorithm:

$$
\rho_{t}^{B A}(a \mid s)=\frac{\rho_{t}^{B A}(a) e^{\gamma_{t} \hat{R}_{t}(a, s)}}{\sum_{a} \rho_{t}^{B A}(a) e^{\gamma_{t} \hat{R}_{t}(a, s)}}, \quad \rho_{t}^{B A}(a)=\frac{1}{N} \sum_{s} \rho_{t}^{B A}(a \mid s)
$$

Running a similar type of iterations in our case would be prohibitively expensive, since they require iteration over all states $s \in \mathcal{S}$ at each round of the game. We approximate these iterations by approximating the marginal distribution over the actions by a running average:

$$
\begin{equation*}
\tilde{\rho}_{t+1}^{e x p}(a)=\frac{N-1}{N} \tilde{\rho}_{t}^{e x p}(a)+\frac{1}{N} \tilde{\rho}_{t}^{e x p}\left(a \mid S_{t}\right) . \tag{7}
\end{equation*}
$$

Since $\rho_{t}^{\text {exp }}(a \mid s)$ is bounded from zero by a decreasing sequence $\varepsilon_{t+1}$, the same automatically holds for $\tilde{\rho}_{t+1}^{\exp }(a)$ (meaning that in Theorem $1 \epsilon_{t}=\varepsilon_{t}$ ). Note that Theorem 1 holds for any choice of $\rho_{t}(a)$, including (7).

We point out an interesting fact: $\rho_{t}^{e x p}(a)$ propagates information between different states, but Theorem 1 also holds for the uniform distribution $\rho(a)=\frac{1}{K}$, which corresponds to application of EXP3 algorithm in each state independently. If these independent multiarmed bandits independently converge to similar strategies, we still get a tighter regret bound. This happens because the corresponding subspace of the hypothesis space is significantly smaller than the total hypothesis space, which enables us to put a higher prior on it (Seldin and Tishby, 2010). Nevertheless, propagation of information between states via the distribution $\rho_{t}^{e x p}(a)$ helps to achieve even faster convergence of the regret, as we can see from the experiments in the next section.

Comparison with state-of-the-art. We are not aware of algorithms for stochastic multiarmed bandits with side information. The best known to us algorithm for adversarial multiarmed bandits with side information is EXP4.P by Beygelzimer et al. (2011). EXP4.P has $O(\sqrt{K t \ln |\mathcal{H}|})$ regret and $O(K|\mathcal{H}|)$ complexity per game round. In our case $|\mathcal{H}|=K^{N}$, which means that EXP4.P would have $O(\sqrt{K t N \ln K})$ regret and $O\left(K^{N+1}\right)$ computational complexity. For hard problems, where all side information has to be used, our regret bound is inferior to the regret bound of Beygelzimer et al. (2011) due to $O\left(t^{2 / 3}\right)$ dependence on the number of game rounds. However, we believe that this can be improved by a more careful analysis of the existing algorithm. For simple problems the dependence of our regret bound on the number of states is significantly better, up to the point that when the side information is irrelevant for the task we can get $O(\sqrt{K \ln N})$ dependence on the number of states versus $O(\sqrt{N \ln K})$ in EXP4.P. For $N \gg K$ this leads to tighter regret bounds for small $t$ even despite the "incorrect" dependence on $t$ of our bound, and if we improve the analysis it will lead to tighter regret bounds for all $t$. As we said it already, our algorithm is able to filter relevant information from large amounts of side information automatically, whereas in EXP4.P the usage of side information has to be restricted externally through the construction of the hypothesis class.


Figure 1: Behavior of: (a) cumulative regret $\Delta(t)$, (b) bound on instantaneous regret $\Delta\left(\tilde{\rho}_{t}^{e x p}\right)$, and (c) the approximation of mutual information $I_{\rho_{t}}^{\exp }(S ; A)$. "Baseline" in the first graph corresponds to playing $N$ independent multiarmed bandits, one in each state. Each line in the graphs corresponds to an average over 10 repetitions of the experiment.

The second important advantage of our algorithm is the exponential improvement of computational complexity. This is achieved by switching from the space of experts to the state-action space in all our calculations.

## 4 Experiments

We present an experiment on synthetic data that illustrates our results. We take $N=100, K=20$, a uniform distribution over states $(p(s)=0.01)$, and consider four settings, with $H\left(A^{h^{*}}\right)=\ln (1)=$ $0, H\left(A^{h^{*}}\right)=\ln (3) \approx 1, H\left(A^{h^{*}}\right)=\ln (7) \approx 2$, and $H\left(A^{h^{*}}\right)=\ln (20) \approx 3$, respectively. In the first case, the same action is the best in all states (and hence $H\left(A^{h^{*}}\right)=0$ for the optimal hypothesis $h^{*}$ ). In the second case, for the first 33 states the best action is number 1, for the next 33 states the best action is number 2, and for the rest third of the states the best action is number 3 (thus, depending on the state, one of the three actions is the "best" and $\left.H\left(A^{h^{*}}\right)=\ln (3)\right)$. In the third case, there are seven groups of 14 states and each group has its own best action. In the last case, there are 20 groups of 5 states and each of $K=20$ actions is the best in exactly one of the 20 groups. For all states, the reward of the best action in a state has Bernoulli distribution with bias 0.6 and the rewards of all other actions in that state have Bernoulli distribution with bias 0.5 . We run the experiment for $T=4,000,000$ rounds and calculate the cumulative regret $\Delta(t)=\sum_{\tau=1}^{t} \Delta\left(\tilde{\rho}_{\tau}^{\text {exp }}\right)$ and instantaneous regret bound given in (4). For computational efficiency, the mutual information $I_{\rho_{t}^{e x p}}(S ; A)$ is approximated by a running average (see supplementary material for details).

As we can see from the graphs (see Figure 1), the algorithm exhibits sublinear cumulative regret (put attention to the axes' scales). Furthermore, for simple problems (with small $H\left(A^{h^{*}}\right)$ ) the regret grows slower than for complex problems. "Baseline" in Figure 1.a shows the performance of an algorithm with the same parameter values that runs $N$ multiarmed bandits, one in each state independently of other states. We see that for all problems except the hardest one our algorithm performs better than the baseline and for the hardest problem it performs almost as good as the baseline. The regret bound in Figure 1.b provides meaningful values for the simplest problem after 1 million rounds (which is on average 500 samples per state-action pair) and after 4 million rounds for all the problems (the graph starts at $t=10,000$ ). Our estimates of the mutual information $I_{\rho_{t} e^{\exp }}(S ; A)$ reflect $H\left(A^{h^{*}}\right)$ for the corresponding problems (for $H\left(A^{h^{*}}\right)=0$ it converges to zero, for $H\left(A^{h^{*}}\right) \approx 1$ it is approximately one, etc.).

## 5 Proof of Theorem 2

The proof of Theorem 2 is based on PAC-Bayes-Bernstein inequality for martingales (Seldin et al., 2011b). Let $\mathrm{KL}(\rho \| \mu)$ denote the KL-divergence between two distributions (Cover and Thomas, 1991). Let $\left\{Z_{1}(h), \ldots, Z_{n}(h): h \in \mathcal{H}\right\}$ be martingale difference sequences indexed by $h$ with respect to the filtration $\sigma\left(\mathcal{U}_{1}\right), \ldots, \sigma\left(\mathcal{U}_{n}\right)$, where $\mathcal{U}_{i}=\left\{Z_{1}(h), \ldots, Z_{i}(h): h \in \mathcal{H}\right\}$ is the subset of martingale difference variables up to index $i$ and $\sigma\left(\mathcal{U}_{i}\right)$ is the $\sigma$-algebra generated by $\mathcal{U}_{i}$. This means that $\mathbb{E}\left[Z_{i}(h) \mid \sigma\left(\mathcal{U}_{i-1}\right)\right]=0$, where $Z_{i}(h)$ may depend on $Z_{j}\left(h^{\prime}\right)$ for all $j<i$ and $h^{\prime} \in \mathcal{H}$. There might also be interdependence between $\left\{Z_{i}(h): h \in \mathcal{H}\right\}$. Let $\hat{M}_{i}(h)=\sum_{j=1}^{i} Z_{j}(h)$ be
the corresponding martingales. Let $V_{i}(h)=\sum_{j=1}^{i} \mathbb{E}\left[Z_{j}(h)^{2} \mid \sigma\left(\mathcal{U}_{j-1}\right)\right]$ be cumulative variances of the martingales $\hat{M}_{i}(h)$. For a distribution $\rho$ over $\mathcal{H}$ define $\hat{M}_{i}(\rho)=\mathbb{E}_{\rho(h)}\left[\hat{M}_{i}(h)\right]$ and $V_{t}(\rho)=$ $\mathbb{E}_{\rho(h)}\left[V_{t}(h)\right]$ as weighted averages of the martingales and their cumulative variances according to a distribution $\rho$.
Theorem 3 (PAC-Bayes-Bernstein Inequality). Assume that $\left|Z_{i}(h)\right| \leq b$ for all $h$ with probability 1. Fix a prior distribution $\mu$ over $\mathcal{H}$. Pick an arbitrary number $c>1$. Then with probability greater than $1-\delta$ over $\mathcal{U}_{n}$, simultaneously for all distributions $\rho$ over $\mathcal{H}$ that satisfy

$$
\sqrt{\frac{\mathrm{KL}(\rho \| \mu)+\ln \frac{2 m}{\delta}}{(e-2) V_{n}(\rho)}} \leq \frac{1}{c b}
$$

we have

$$
\left|\hat{M}_{n}(\rho)\right| \leq(1+c) \sqrt{(e-2) V_{n}(\rho)\left(\mathrm{KL}(\rho \| \mu)+\ln \frac{2 m}{\delta}\right)}
$$

where $m=\ln \left(\sqrt{\frac{(e-2) n}{\ln \frac{2}{\delta}}}\right) / \ln (c)$, and for all other $\rho$

$$
\left|\hat{M}_{n}(\rho)\right| \leq 2 b\left(K L(\rho \| \mu)+\ln \frac{2 m}{\delta}\right)
$$

Note that $M_{t}(h)=t\left(\Delta(h)-\hat{\Delta}_{t}(h)\right)$ are martingales and their cumulative variances are $V_{t}(h)=$ $\sum_{\tau=1}^{t} \mathbb{E}\left[\left(\left[R_{\tau}^{h^{*}\left(S_{\tau}\right), S_{\tau}}-R_{\tau}^{h\left(S_{\tau}\right), S_{\tau}}\right]-\left[R\left(h^{*}\right)-R(h)\right]\right)^{2} \mid \mathcal{T}_{\tau-1}\right]$. In order to apply Theorem 3 we have to derive an upper bound on $V_{t}\left(\rho_{t}^{\text {exp }}\right),{ }^{1}$ a prior $\mu(h)$ over $\mathcal{H}$, and calculate (or upper bound) the $\operatorname{KL}$-divergence $\operatorname{KL}\left(\rho_{t}^{e x p} \| \mu\right)$. This is done in the following three lemmas.
Lemma 3. If $\left\{\varepsilon_{1}, \varepsilon_{2}, \ldots\right\}$ is a decreasing sequence, such that $\varepsilon_{t} \leq \min _{a, s} \pi_{t}(a \mid s)$, then for all $h$ :

$$
V_{t}(h) \leq \frac{2 t}{\varepsilon_{t}}
$$

The proof of the lemma is provided in the supplementary material. Lemma 3 provides an immediate, but crude, uniform upper bound on $V_{t}(h)$, which yields $V_{t}\left(\rho_{t}^{\text {exp }}\right) \leq \frac{2 t}{\varepsilon_{t}}$. Since our algorithm concentrates on $h$-s with small $\Delta(h)$, which, in turn, concentrate on the best action in each state, the variance $V_{t}(h)$ for the corresponding $h$-s is expected to be of the order of $2 K t$ and not $\frac{2 t}{\varepsilon_{t}}$. However, we were not able to prove yet that the probability $\rho_{t}^{e x p}(h)$ of the remaining hypotheses (those with large $\Delta(h)$ ) gets sufficiently small (of order $K \varepsilon_{t}$ ), so that the weighted cumulative variance would be of order 2 Kt . Nevertheless, this seems to hold in practice starting from relatively small values of $t$ (Seldin et al., 2011a). Improving the upper bound on $V_{t}\left(\rho_{t}^{\text {exp }}\right)$ will improve the regret bound, but for the moment we present the regret bound based on the crude upper bound $V_{t}\left(\rho_{t}^{\text {exp }}\right) \leq \frac{2 t}{\varepsilon_{t}}$.
The remaining two lemmas, which define a prior $\mu$ over $\mathcal{H}$ and bound $\operatorname{KL}(\rho \| \mu)$, are due to Seldin and Tishby (2010).
Lemma 4. It is possible to define a distribution $\mu$ over $\mathcal{H}$ that satisfies:

$$
\begin{equation*}
\mu(h) \geq e^{-N H\left(A^{h}\right)-K \ln N-K \ln K} . \tag{8}
\end{equation*}
$$

Lemma 5. For the distribution $\mu$ that satisfies (8) and any distribution $\rho(a \mid s)$ :

$$
\mathrm{KL}(\rho \| \mu) \leq N I_{\rho}(S ; A)+K \ln N+K \ln K
$$

Substitution of the upper bounds on $V_{t}\left(\rho_{t}^{e x p}\right)$ and $\mathrm{KL}\left(\rho_{t}^{e x p} \| \mu\right)$ into Theorem 3 yields Theorem 2.

## 6 Discussion

We presented PAC-Bayesian analysis of stochastic multiarmed bandits with side information. Our analysis provides data-dependent algorithm and data-dependent regret analysis for this problem. The selection of task-relevant side information is delegated from the user to the algorithm. We also provide a general framework for deriving data-dependent algorithms and analyses for many other stochastic problems with limited feedback. The analysis of the variance of our algorithm still waits to be improved and will be addressed in future work.

[^0]
## Acknowledgments

We would like to thank all the people with whom we discussed this work and, in particular, Nicolò Cesa-Bianchi, Gábor Bartók, Elad Hazan, Csaba Szepesvári, Miroslav Dudík, Robert Shapire, John Langford, and the anonymous reviewers, whose comments helped us to improve the final version of this manuscript. This work was supported in part by the IST Programme of the European Community, under the PASCAL2 Network of Excellence, IST-2007-216886, and by the European Community's Seventh Framework Programme (FP7/2007-2013), under grant agreement $N^{\circ} 231495$. This publication only reflects the authors' views.

## References

Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. Machine Learning, 47, 2002a.

Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. SIAM Journal of Computing, 32(1), 2002b.
Arindam Banerjee. On Bayesian bounds. In Proceedings of the International Conference on Machine Learning (ICML), 2006.
Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert Schapire. Contextual bandit algorithms with supervised learning guarantees. In Proceedings on the International Conference on Artificial Intelligence and Statistics (AISTATS), 2011.
Thomas M. Cover and Joy A. Thomas. Elements of Information Theory. John Wiley \& Sons, 1991.
Leslie Pack Kaelbling. Associative reinforcement learning: Functions in $k$-DNF. Machine Learning, 15, 1994.
John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In Advances in Neural Information Processing Systems (NIPS), 2007.
David McAllester. Some PAC-Bayesian theorems. In Proceedings of the International Conference on Computational Learning Theory (COLT), 1998.
Matthias Seeger. PAC-Bayesian generalization error bounds for Gaussian process classification. Journal of Machine Learning Research, 2002.
Yevgeny Seldin and Naftali Tishby. PAC-Bayesian analysis of co-clustering and beyond. Journal of Machine Learning Research, 11, 2010.
Yevgeny Seldin, Nicolò Cesa-Bianchi, Peter Auer, François Laviolette, and John Shawe-Taylor. PAC-BayesBernstein inequality for martingales and its application to multiarmed bandits. 2011a. In review. Preprint available at http://arxiv.org/abs/1110.6755.
Yevgeny Seldin, François Laviolette, Nicolò Cesa-Bianchi, John Shawe-Taylor, and Peter Auer. PAC-Bayesian inequalities for martingales. 2011b. In review. Preprint available at http://arxiv.org/abs/1110.6886.

John Shawe-Taylor and Robert C. Williamson. A PAC analysis of a Bayesian estimator. In Proceedings of the International Conference on Computational Learning Theory (COLT), 1997.
John Shawe-Taylor, Peter L. Bartlett, Robert C. Williamson, and Martin Anthony. Structural risk minimization over data-dependent hierarchies. IEEE Transactions on Information Theory, 44(5), 1998.
Alexander L. Strehl, Chris Mesterharm, Michael L. Littman, and Haym Hirsh. Experience-efficient learning in associative bandit problems. In Proceedings of the International Conference on Machine Learning (ICML), 2006.

Richard S. Sutton and Andrew G. Barto. Reinforcement Learning: An Introduction. MIT Press, 1998.
Naftali Tishby and Daniel Polani. Information theory of decisions and actions. In Vassilis Cutsuridis, Amir Hussain, John G. Taylor, and Daniel Polani, editors, Perception-Reason-Action Cycle: Models, Algorithms and Systems. Springer, 2010.
Naftali Tishby, Fernando Pereira, and William Bialek. The information bottleneck method. In Allerton Conference on Communication, Control and Computation, 1999.
Leslie G. Valiant. A theory of the learnable. Communications of the Association for Computing Machinery, 27 (11), 1984.


[^0]:    ${ }^{1}$ Seldin et al. (2011b) show that $V_{n}(\rho)$ can be replaced by an upper bound everywhere in Theorem 3.

