# COMS6998-11: Homework 1 Solutions 

Akshay Krishnamurthy<br>akshay@cs.umass.edu

1. Let us start by setting the RHS of Bernstein's inequality to be at most $\delta$ and re-arranging:

$$
\exp \left(-\frac{t^{2} / 2}{\sum_{i=1}^{n} \mathbb{E}\left[X_{i}^{2}\right]+M t / 3}\right) \leq \delta \Leftrightarrow t^{2} \geq 2\left(\sum_{i=1}^{n} \mathbb{E}\left[X_{i}^{2}\right]\right) \log (1 / \delta)+\frac{2 M t}{3} \log (1 / \delta)
$$

This is a quadratic inequality in $t$ and we can find the roots using the quadratic formula. That is somewhat tedious, so a shortcut is to decompose $t=t_{1}+t_{2}$ where $t_{1}$ and $t_{2}$ are chosen to ensure different parts of this inequality hold. Specifically we set

$$
t_{1}=\sqrt{2\left(\sum_{i=1}^{n} \mathbb{E}\left[X_{i}^{2}\right]\right) \log (1 / \delta)}, \quad \text { and } \quad t_{2}=\frac{2 M}{3} \log (1 / \delta)
$$

Now, we can check

$$
t^{2}=t_{1}^{2}+2 t_{1} t_{2}+t_{2}^{2} \geq t_{1}^{2}+t_{2} t
$$

which, using the definitions of $t_{1}$ and $t_{2}$ is precisely the inequality we want to satisfy, after dividing by $n$ to account for the difference between $\bar{X}$ and $\sum_{i=1}^{n} X_{i}$.
For the generalization bound, fix a function $f \in \mathcal{F}$ and define random variables $Z_{1}, \ldots, Z_{n}$ where $Z_{i}=$ $\mathbf{1}\left\{f\left(x_{i}\right) \neq y_{i}\right\}$. Observe that $\hat{R}_{n}(f)=\frac{1}{n} \sum_{i=1}^{n} Z_{i}$ and $R(f)=\mathbb{E}\left[Z_{i}\right]$. So we will apply Bernstein's inequality on the random variables

$$
\hat{R}_{n}(f)-R(f)=\frac{1}{n} \sum_{i=1}^{n}\left(\mathbf{1}\left\{f\left(x_{i}\right) \neq y_{i}\right\}-R(f)\right)
$$

A critical step in the argument is that the variance term that appears in Bernstein's inequaltiy can be upper bounded by the mean:

$$
\mathbb{E}\left[\left(\mathbf{1}\left\{f\left(x_{i}\right) \neq y_{i}\right\}-R(f)\right)^{2}\right] \leq \mathbb{E}\left[\mathbf{1}\left\{f\left(x_{i}\right) \neq y_{i}\right\}^{2}\right] \leq \mathbb{E}\left[\mathbf{1}\left\{f\left(x_{i}\right) \neq y_{i}\right\}\right]=R(f)
$$

Then, by our variant of Bernstein's inequality and a union bound, we have that with probability at least $1-\delta$ :

$$
\forall f \in \mathcal{F}:\left|\hat{R}_{n}(f)-R(f)\right| \leq \sqrt{\frac{2 R(f) \log (2|\mathcal{F}| / \delta)}{n}}+\frac{2 \log (2|\mathcal{F}| / \delta)}{3}
$$

Using this with $\hat{f}_{n}$ along with the fact that $\hat{f}_{n}$ is the ERM yields

$$
\begin{aligned}
R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right) & \leq \hat{R}_{n}\left(\hat{f}_{n}\right)-\hat{R}_{n}\left(f^{\star}\right)+\sqrt{\frac{2 R\left(\hat{f}_{n}\right) \log (2|\mathcal{F}| / \delta)}{n}}+\sqrt{\frac{2 R\left(f^{\star}\right) \log (2|\mathcal{F}| / \delta)}{n}}+\frac{4 \log (2|\mathcal{F}| / \delta)}{3} \\
& \leq \sqrt{\frac{2 R\left(\hat{f}_{n}\right) \log (2|\mathcal{F}| / \delta)}{n}}+\sqrt{\frac{2 R\left(f^{\star}\right) \log (2|\mathcal{F}| / \delta)}{n}}+\frac{4 \log (2|\mathcal{F}| / \delta)}{3}
\end{aligned}
$$

The only problem is the first term in the above bound since it depends on $R\left(\hat{f}_{n}\right)$ rather than $R\left(f^{\star}\right)$. On the other hand, the bound itself is showing that $R\left(\hat{f}_{n}\right) \approx R\left(f^{\star}\right)$ so intuitively we should be able to swap the $R\left(\hat{f}_{n}\right)$ term on the RHS with $R\left(f^{\star}\right)$ without paying too much. Indeed

$$
\begin{aligned}
\sqrt{\frac{2 R\left(\hat{f}_{n}\right) \log (2|\mathcal{F}| / \delta)}{n}} & =\sqrt{\frac{2\left(R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right)+R\left(f^{\star}\right)\right) \log (2|\mathcal{F}| / \delta)}{n}} \\
& \leq \sqrt{\frac{2\left(R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right)\right) \log (2|\mathcal{F}| / \delta)}{n}}+\sqrt{\frac{2 R\left(f^{\star}\right) \log (2|\mathcal{F}| / \delta)}{n}} \\
& \leq \frac{R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right)}{2}+\frac{\log (2|\mathcal{F}| / \delta)}{n}+\sqrt{\frac{2 R\left(f^{\star}\right) \log (2|\mathcal{F}| / \delta)}{n}}
\end{aligned}
$$

Here the first inequality is that $\sqrt{a+b} \leq \sqrt{a}+\sqrt{b}$ while the second is that $\sqrt{a b} \leq a / 2+b / 2$. Plugging this bound in above, we have

$$
R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right) \leq \frac{R\left(\hat{f}_{n}\right)-R\left(f^{\star}\right)}{2}+2 \sqrt{\frac{2 R\left(f^{\star}\right) \log (2|\mathcal{F}| / \delta)}{n}}+\frac{7 \log (2|\mathcal{F}| / \delta)}{3 n}
$$

By bringing the first term from the RHS to the LHS and then multiplying by 2, we obtain the result.
2. The first observation is that for any epoch $i$ if an arm survives into epoch $i$ then at the end of the epoch it has been played $N_{i}=\sum_{j=1}^{i} 2^{j} \geq 2^{i}$. In particular this means that all surviving arms have been played the same number of times.
Let $a^{\star}=\operatorname{argmax}_{a} \mu(a)$ denote the best arm. Next, we show that $a^{\star}$ is never eliminated. To do this, recall that we set the confidence interval such that

$$
\forall t \in[T], a \in \mathcal{A}:\left|\hat{\mu}_{t}(a)-\mu(a)\right| \leq \operatorname{conf}_{t}(a)
$$

We show that for any epoch $i, a^{\star}$ is not eliminated. This follows since for any $a \neq a^{\star}$

$$
\hat{\mu}\left(a^{\star}\right)+\operatorname{conf}(a) \geq \mu\left(a^{\star}\right) \geq \mu(a) \geq \hat{\mu}(a)-\operatorname{conf}(a)
$$

where the first and third inequality are by the property of the confidence interval and the second inequality is by the optimality of $a^{\star}$.
Next we claim that any arm that survives an elimination step cannot be too suboptimal. Indeed
$\mu(a) \geq \hat{\mu}(a)-\operatorname{conf}(a)=\hat{\mu}(a)+\operatorname{conf}(a)-2 \operatorname{conf}(a) \geq \hat{\mu}\left(a^{\star}\right)-\operatorname{conf}\left(a^{\star}\right)-2 \operatorname{conf}(a) \geq \mu\left(a^{\star}\right)-2 \operatorname{conf}\left(a^{\star}\right)-2 \operatorname{conf}(a)$
Since the confidence interval for all surviving arms is the same, we see that if arm $a$ survives through epoch $i$ then its suboptimality is roughly $\sqrt{1 / 2^{i}}$.
For the regret analysis, the bookkeeping has to be done somewhat carefully. We decompose the regret across epochs. In each epoch the regret we incur depends on the number of active arms $\left|\mathcal{A}_{i}\right|$, the epoch length, and the confidence interval. Formally the regret in epoch $i$ is bounded by $4\left|\mathcal{A}_{i}\right| \cdot 2^{i} \cdot \operatorname{conf}_{i-1}$ (below we will not track the constants precisely). Let $I$ denote the index of the last complete epoch and let $I_{a}$ denote the index of the epoch when $a$ is eliminated (or $I$ if $a$ is never eliminated) and let $N_{a}$ denote the number of plays for arm $a$ up to and including epoch $I$.
First we claim that we do not incur too much regret in the final epoch $I+1$, which does not complete. Suppose that $\left|\mathcal{A}_{I+1}\right|=k>0$. Then observe that $I \leq\left\lfloor\log _{2}(T / K)\right\rfloor$, since epoch $I$ lasts for at least $k 2^{I}$ rounds. The regret in epoch $I+1$ can be crudely upper bounded by

$$
k \cdot 2^{I+1} \cdot \operatorname{conf}_{I} \leq c k 2^{I+1} \cdot \sqrt{\frac{\log (A T / \delta)}{2^{I}}} \leq c k \sqrt{2^{I} \log (A T / \delta)} \leq c \sqrt{k T \log (A T / \delta)} \leq O(\sqrt{A T \log (A T / \delta)})
$$

For the other epochs we have

$$
\begin{aligned}
\sum_{i=1}^{I}\left|\mathcal{A}_{i}\right| \cdot 2^{i} \cdot \operatorname{conf}_{i-1} & \leq c \sum_{i=1}^{I}\left|\mathcal{A}_{i}\right| \sqrt{2^{i} \cdot \log (A T / \delta)} \\
& \leq c \sum_{a} \sum_{i=1}^{I_{a}} \sqrt{2^{i} \cdot \log (A T / \delta)} \\
& \leq c \sum_{a} \sqrt{N_{a} \cdot \log (A T / \delta)} \\
& \leq c \sqrt{A T \log (A T / \delta)}
\end{aligned}
$$

To pass to the second line, we simply re-arrange the sum so that for each arm $a$ we consider all the epochs until $a$ is eliminated. Then we bound this geometric series (the constant $c$ changes from line to line). Finally we use Cauchy-Schwarz since $\sum_{a} N_{a} \leq T$.
3. For item (a) our choice of $\beta$ ensures that $\left\|\hat{\theta}_{t}-\theta^{\star}\right\|_{\Lambda_{t}} \leq \beta$, just as we saw in class. Now considering $a^{\star}$ and some other action $b$, we have

$$
\begin{aligned}
\left\langle\hat{\theta}_{t}, x_{b}-x_{a^{\star}}\right\rangle & =\left\langle\hat{\theta}_{t}-\theta^{\star}, x_{b}-x_{a^{\star}}\right\rangle+\left\langle\theta^{\star}, x_{b}-x_{a^{\star}}\right\rangle \\
& \leq\left\langle\hat{\theta}_{t}-\theta^{\star}, x_{b}-x_{a^{\star}}\right\rangle \\
& \leq\left\|\hat{\theta}_{t}-\theta^{\star}\right\|_{\Lambda_{t}} \cdot\left\|x_{b}-x_{a^{\star}}\right\|_{\Lambda_{t}^{-1}} \\
& \leq \beta\left\|x_{b}-x_{a^{\star}}\right\|_{\Lambda_{t}^{-1}}
\end{aligned}
$$

Here the first inequality uses the optimality of $a^{\star}$ (for the true parameter $\theta^{\star}$ ). The second is Cauchy-Schwarz and the third uses our guarantee on $\hat{\theta}_{t}$. This proves that $a^{\star}$ is never eliminated.
For item (b), the total regret across all rounds is

$$
\sum_{t=1}^{T}\left\langle\theta^{\star}, x_{a^{\star}}-x_{a_{t}}\right\rangle \leq \sum_{t=1}^{T} \sum_{a} p_{t}(a)\left\langle\theta^{\star}, x_{a^{\star}}-x_{a}\right\rangle+O(\sqrt{T \log (1 / \delta)})
$$

where the latter holds with probability $1-\delta$ by an application of Azuma's inequality. Now let us focus on one of the terms in the first sum:

$$
\begin{aligned}
\sum_{a} p_{t}(a)\left\langle\theta^{\star}, x_{a^{\star}}-x_{a}\right\rangle & =\left\langle\theta^{\star}, x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\rangle \\
& =\left\langle\theta^{\star}-\hat{\theta}_{t}, x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\rangle+\left\langle\hat{\theta}_{t}, x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\rangle \\
& \leq \beta \cdot\left\|x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\|_{\Lambda_{t}^{-1}}+\left\langle\hat{\theta}_{t}, x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\rangle
\end{aligned}
$$

The first term is at a point where we can apply Eq. (4), so let us continue with the third term. By applying Eq. (3) and then the triangle inequality we obtain

$$
\left\langle\hat{\theta}_{t}, x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\rangle \leq \beta \sum_{a} p_{t}(a)\left\|x_{a^{\star}}-x_{a}\right\|_{\Lambda_{t}^{-1}} \leq \beta\left\|x_{a^{\star}}-\mathbb{E}_{a \sim p_{t}}\left[x_{a}\right]\right\|_{\Lambda_{t}^{-1}}+\beta \sum_{a} p_{t}(a)\left\|\mathbb{E}_{b \sim p_{t}}\left[x_{b}\right]-x_{a}\right\|_{\Lambda_{t}^{-1}}
$$

All of these terms are bounded by Eq. (4). Together, we obtain the bound

$$
\sum_{a} p_{t}(a)\left\langle\theta^{\star}, x_{a^{\star}}-x_{a}\right\rangle \leq 3 \beta \sqrt{\operatorname{tr}\left(\Lambda_{t}^{-1} \cdot M_{t}\right)}
$$

For item (c) after the application of Azuma's inequality from above, and using our answer to part (b) the regret is bounded as

$$
\sum_{t=1}^{T}\left\langle\theta^{\star}, x_{a^{\star}}-x_{a_{t}}\right\rangle \leq 3 \beta \sum_{t=1}^{T} \sqrt{\operatorname{tr}\left(\Lambda_{t}^{-1} \cdot M_{t}\right)}+O(\sqrt{T \log (1 / \delta)})
$$

We will focus on the first term. By Cauchy-Schwarz and then by using the concentration inequality provided, we bound this term as

$$
\begin{aligned}
\sum_{t=1}^{T} \sqrt{\operatorname{tr}\left(\Lambda_{t}^{-1} \cdot M_{t}\right)} & \leq \sqrt{T} \cdot \sqrt{\sum_{t=1}^{T} \operatorname{tr}\left(\Lambda_{t}^{-1} \cdot M_{t}\right)} \\
& \leq \sqrt{T} \cdot \sqrt{\sum_{t=1}^{T} \operatorname{tr}\left(\Lambda_{t}^{-1} x_{a_{t}} x_{a_{t}}^{\top}\right)+8 \log (2 / \delta)} \\
& \leq \sqrt{T} \cdot \sqrt{\sum_{t=1}^{T} x_{a_{t}}^{\top} \Lambda_{t}^{-1} x_{a_{t}}}+O(\sqrt{T \log (1 / \delta)})
\end{aligned}
$$

Now we can apply the elliptical potential lemma, since $\Lambda_{t+1}=\Lambda_{t}+x_{a_{t}} x_{a_{t}}^{\top}$. The only minor difference is that we don't have the $\min (1, \cdot)$ in our sum. However since we started with $\Lambda_{1}=I$, this does not affect the argument too much. In particular the only change is in the first step, where we use the Woodbury identity, Here we get

$$
x_{a_{t}}^{\top} \Lambda_{t+1}^{-1} x_{a_{t}}=\frac{x_{a_{t}^{\top}} \Lambda_{t}^{-1} x_{a_{t}}}{1+x_{a_{t}^{\top}} \Lambda_{t}^{-1} x_{a_{t}}} \geq\left(1+B^{2}\right)^{-1} x_{a_{t}^{\top}} \Lambda_{t}^{-1} x_{a_{t}}
$$

which follows since $\Lambda_{t}^{-1} \preceq I$ and we assume that the features are bounded as $\left\|x_{a_{t}}\right\|_{2} \leq B$. Using the rest of the elliptical potential as is, we can conclude that

$$
\sum_{t=1}^{T} x_{a_{t}}^{\top} \Lambda_{t}^{-1} x_{a_{t}} \leq \frac{1}{1+B^{2}} d \log \left(1+T B^{2} / d\right)
$$

and using this yields the desired regret bound.
4. For part (a) observe that

$$
\begin{aligned}
& \mathbb{E}\left[\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-r_{t}\right)^{2}-\left(f^{\star}\left(x_{t}, a_{t}\right)-r_{t}\right)^{2}\right] \\
& =\mathbb{E}\left[\hat{y}_{t}\left(x_{t}, a_{t}\right)^{2}-f^{\star}\left(x_{t}, a_{t}\right)^{2}-2 r_{t}\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\right] \\
& =\mathbb{E}\left[\hat{y}_{t}\left(x_{t}, a_{t}\right)^{2}-f^{\star}\left(x_{t}, a_{t}\right)^{2}-2 \mu_{t}\left(x_{t}, a_{t}\right)\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\right] \\
& =\mathbb{E}\left[\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)^{2}-2\left(\mu_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\right]
\end{aligned}
$$

Therefore

$$
\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{T}\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)^{2}\right] & \leq \mathbb{E}\left[\operatorname{Reg}_{\mathrm{sq}}(T)\right]+\mathbb{E}\left[\sum_{t=1}^{T} 2\left(\mu_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)\right] \\
& \leq \mathbb{E}\left[\operatorname{Reg}_{\mathrm{sq}}(T)\right]+\mathbb{E}\left[\sum_{t=1}^{T} 2\left(\mu_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)^{2}+\sum_{t=1}^{T} \frac{1}{2}\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)^{2}\right] \\
& \leq \mathbb{E}\left[\operatorname{Reg}_{\mathrm{sq}}(T)\right]+2 \varepsilon^{2} T+\mathbb{E}\left[\sum_{t=1}^{T} \frac{1}{2}\left(\hat{y}_{t}\left(x_{t}, a_{t}\right)-f^{\star}\left(x_{t}, a_{t}\right)\right)^{2}\right]
\end{aligned}
$$

Re-arranging this proves the result.

For the part (b) we need to prove a "per-round" inequality, but first let us see what the per-round inequality should be. The first thing to do is to move from the true reward function $\mu$ to the function $f^{\star}$ that is in our function class. For this, observe that for any $x_{t}$ and $a_{t}$

$$
\begin{aligned}
\max _{a} \mu\left(x_{t}, a\right)-\mu\left(x_{t}, a_{t}\right) & =\max _{a} \mu\left(x_{t}, a\right)-\max _{a} f^{\star}\left(x_{t}, a\right)+\max _{a} f^{\star}\left(x_{t}, a\right)-f^{\star}\left(x_{t}, a_{t}\right)+f^{\star}\left(x_{t}, a_{t}\right)-\mu\left(x_{t}, a_{t}\right) \\
& \leq 2 \varepsilon+\max _{a} f^{\star}\left(x_{t}, a\right)-f^{\star}\left(x_{t}, a_{t}\right),
\end{aligned}
$$

where the inequality holds by our assumption on $f^{\star}$ and $\mu$ (and a similar proof to what we saw in the contraction lemma to relate the two max's).
So from now on we can work with $\mathbb{E}_{a_{t} \sim p_{t}}\left[\max _{a} f^{\star}\left(x_{t}, a\right)-f^{\star}\left(x_{t}, a_{t}\right)\right]$. Let us condense the notation and drop dependence on $x_{t}$ and $t$ altogether. Let $a^{\star}=\operatorname{argmax}_{a} f^{\star}(a)$ and let $y$ denote the vector of our predictions and let $b=\operatorname{argmax}_{a} y(a)$. So now we want to bound

$$
\mathbb{E}_{a \sim p}\left[f^{\star}\left(a^{\star}\right)-f^{\star}(a)\right]-\frac{\gamma}{4} \mathbb{E}_{a \sim p}\left[\left(y(a)-f^{\star}(a)\right)^{2}\right]
$$

and exactly our analysis for SquareCB shows that this is $O(A / \gamma)$. Putting things together, we have

$$
\begin{aligned}
\operatorname{Regret}(T) & :=\sum_{t=1}^{T} \max _{a} \mu\left(x_{t}, a\right)-\mathbb{E}_{a_{t} \sim p_{t}}\left[\mu\left(x_{t}, a_{t}\right)\right] \\
& \leq 2 T \varepsilon+\sum_{t=1}^{T} \max _{a} f^{\star}\left(x_{t}, a\right)-\mathbb{E}_{a_{t} \sim p_{t}} f^{\star}\left(x_{t}, a_{t}\right) \\
& \leq 2 T \varepsilon+\frac{A T}{\gamma}+\frac{\gamma}{4} \sum_{t=1}^{T} \mathbb{E}_{a \sim p_{t}}\left[\left(\hat{y}_{t}\left(x_{t}, a\right)-f^{\star}\left(x_{t}, a\right)\right)^{2}\right] \\
& \lesssim 2 T \varepsilon+\frac{A T}{\gamma}+\frac{\gamma}{4}\left(2 \operatorname{Reg}_{\mathrm{sq}}(T)+2 \varepsilon^{2} T\right)
\end{aligned}
$$

Here the $\lesssim$ notation hides a concentration argument to relate $\operatorname{Reg}_{\text {sq }}(T)$ to its expectation. Now if we optimize for $\gamma$ we get $\gamma=\Theta\left(\sqrt{\frac{A T}{\operatorname{Reg}_{\mathrm{sq}}(T)+2 \varepsilon^{2} T}}\right.$ and plugging this in yields

$$
O\left(T \varepsilon+\sqrt{A T\left(\operatorname{Reg}_{\mathrm{sq}}(T)+2 \varepsilon^{2} T\right)}\right) \leq O\left(\varepsilon T \sqrt{A}+\sqrt{A T \operatorname{Reg}_{\mathrm{sq}}(T)}\right)
$$

