Mehryar Mohri
Advanced Machine Learning 2024
Courant Institute of Mathematical Sciences
Homework assignment 1
February 20, 2024
Due: March 05, 2024

1 Short proof of Hoeffding's lemma

Let X be a random variable with $\mathbb{E}[X] = 0$ and $|X| \leq 1$. Show that for any $t > 0$, $\mathbb{E}[e^{tX}] \leq e^{t^2/2}$. *Hint*: use the convexity of exponential to derive $\mathbb{E}[e^{tX}] \leq \cosh(t)$.

Solution: By convexity of the exponential function, for any $t > 0$,

$$e^{tX} \leq \frac{1-X}{2} e^{-t} + \frac{1+X}{2} e^{t},$$

thus, using $(2n)! = (2n)(2n-1)\cdots \geq (2n)(2(n-1))\cdots = 2^n n!$, we have

$$\mathbb{E}[e^{tX}] \leq \frac{1}{2}(e^{-t} + e^{t}) = \cosh(t) = \sum_{n=0}^{+\infty} \frac{t^{2n}}{(2n)!} \leq \sum_{n=0}^{+\infty} \frac{t^{2n}}{2^n n!} = e^{t^2/2}.$$

□

2 Small loss bound

In lecture 3, we used the bound on the regret of Randomized Weighted Majority (RWM): $R_T \leq O(\sqrt{L_T^{\min} \log N})$, assuming $L_T^{\min} \neq 0$. Here, you are asked to give a proof. Using the proof given in Lecture 1 for the regret of RWM and the inequality $\mathcal{L}_T \leq \frac{\log N}{1-\beta} + (2-\beta)\mathcal{L}_T^{\min}$, show that for a suitable choice of β, we have $R_T \leq 4\sqrt{\mathcal{L}_T^{\min} \log N}$.

Solution: In view of the inequality, we can write

$$R_T \leq \frac{\log N}{1-\beta} + (1-\beta)\mathcal{L}_T^{\min},$$

for $\beta \in [1/2, 1)$, that is $(1-\beta) \in (0, 1/2]$. The right-hand side is a convex function of $(1-\beta)$. Assuming $\mathcal{L}_T^{\min} \neq 0$, the minimizer over $\mathbb{R}_+$ is $\sqrt{\frac{\log N}{\mathcal{L}_T^{\min}}}$, thus the minimizer over $(0, 1/2]$ is achieved for $1 - \beta = \min\left\{\frac{1}{2}, \sqrt{\frac{\log N}{\mathcal{L}_T^{\min}}}\right\}$. Plugging in this value gives: $R_T \leq 4\sqrt{\mathcal{L}_T^{\min} \log N}$. □

3 Weighted online-to-batch

Let ℓ be a loss function convex with respect to its first argument and bounded by one. Let $h_1, \ldots, h_T$ be the hypotheses returned by an on-line learning algorithm $\mathcal{A}$ with regret R_T when sequentially processing $(x_t, y_t)_{t=1}^T$, drawn i.i.d. according to some distribution $\mathcal{D}$.

1. Fix some arbitrary non-negative weights $q_1, \ldots, q_T$ summing to one. Then, show that with probability at least $1-\delta$, the hypothesis $h = \sum_{t=1}^T q_t h_t$ satisfies each of the following inequalities:

$$\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] \leq \sum_{t=1}^T q_t \ell(h_t(x_t), y_t) + \|q\|_2 \sqrt{2\log(1/\delta)}$$

$$\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] \leq \inf_{h\in\mathcal{H}} \mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] + \frac{R_T}{T} + \|q - u\|_1 + 2\|q\|_2 \sqrt{2\log(1/\delta)},$$

where q is the vector with components q_t and u the uniform vector with all components equal to $1/T$.

Solution: Let $R(h_t) = \mathbb{E}_{(x,y)\sim\mathcal{D}}[\ell(h_t(x), y)]$. Let Z_t be the random variable defined by $Z_t = q_t\ell(h_t(x_t), y_t) - q_t R(h_t)$ and let $\mathcal{F}_t$ denote the filtration associated to the sample process. We have $|Z_t| \leq q_t$ and $\mathbb{E}[Z_t \mid \mathcal{F}_{t-1}] = \mathbb{E}[q_t\ell(h_t(x_t), y_t) \mid h_t] - q_t R(h_t) = q_t R(h_t) - q_t R(h_t) = 0$. Thus, by Azuma's inequality, for any $\delta > 0$, with probability at least $1 - \delta$, the following holds:

$$\sum_{t=1}^T q_t R(h_t) \leq \sum_{t=1}^T q_t \ell(h_t(x_t), y_t) + \|q\|_2 \sqrt{2\log(1/\delta)} \quad (1)$$

$$\sum_{t=1}^T q_t \ell(h_t(x_t), y_t) \leq \sum_{t=1}^T q_t R(h_t) + \|q\|_2 \sqrt{2\log(1/\delta)} \quad (2)$$

Since ℓ is convex with respect to its first argument, by Jensen's inequality, we have $\mathbb{E}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] = \mathbb{E}_{(x,y)\sim\mathcal{D}}[\ell(\sum_{t=1}^T q_t h_t(x), y)] \leq \sum_{t=1}^T q_t R(h_t)$. Thus, by (1), for any $\delta > 0$, the following holds with probability at least $1 - \delta$,

$$\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] \leq \sum_{t=1}^T q_t \ell(h_t(x_t), y_t) + \|q\|_2 \sqrt{2\log(1/\delta)}. \quad (3)$$

This proves the first statement. To prove the second claim, we will bound the empirical error in terms of the regret. For any $h^* \in \mathcal{H}$, we can write using $\inf_{h\in\mathcal{H}} \frac{1}{T}\sum_{t=1}^T \ell(h(x_t), y_t) \leq \frac{1}{T}\sum_{t=1}^T \ell(h^*(x_t), y_t)$:

$$\begin{aligned}
&\sum_{t=1}^T q_t \ell(h_t(x_t), y_t) - \sum_{t=1}^T q_t \ell(h^*(x_t), y_t) \\
&= \sum_{t=1}^T \left(q_t - \frac{1}{T}\right)[\ell(h_t(x_t), y_t) - \ell(h^*(x_t), y_t)] + \frac{1}{T}\sum_{t=1}^T [\ell(h_t(x_t), y_t) - \ell(h^*(x_t), y_t)] \\
&\leq \|q - u\|_1 + \frac{1}{T}\sum_{t=1}^T \ell(h_t(x_t), y_t) - \inf_{h\in\mathcal{H}} \frac{1}{T}\sum_{t=1}^T \ell(h(x_t), y_t) \\
&\leq \|q - u\|_1 + \frac{R_T}{T}
\end{aligned}$$

Now, by definition of the infimum, for any $\epsilon > 0$, there exists $h^* \in \mathcal{H}$ such that $\mathbb{E}_{(x,y)\sim\mathcal{D}}[\ell(h^*(x), y)] \leq \inf_{h\in\mathcal{H}} \mathbb{E}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] + \epsilon$. For that choice of h^*, in view of (3), with probability at least $1 - \delta/2$, the following holds:

$$\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] \leq \sum_{t=1}^T q_t \ell(h^*(x_t), y_t) + \|q - u\|_1 + \frac{R_T}{T} + \|q\|_2 \sqrt{2\log(1/\delta)}.$$

By (2), for any $\delta > 0$, with probability at least $1 - \delta/2$,

$$\sum_{t=1}^T q_t \ell(h^*(x_t), y_t) \leq \mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h^*(x), y)] + \|q\|_2 \sqrt{2\log(1/\delta)}.$$

Combining these last two inequalities, by the union bound, with probability at least $1 - \delta$, the following holds:

$$\begin{aligned}
\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] \leq \inf_{h\in\mathcal{H}} &\mathop{\mathbb{E}}_{(x,y)\sim\mathcal{D}}[\ell(h(x), y)] + \epsilon + \frac{R_T}{T} \\
&+ \|q - u\|_1 + 2\|q\|_2 \sqrt{2\log(1/\delta)}.
\end{aligned}$$

Since this inequality holds for all $\epsilon > 0$, it implies the second statement. □

2. Here, we seek to prove a bound that holds uniformly for all weight vectors q in some set. To do so, we consider a weight vector p that serves as a *reference*. A natural reference in this context could be for example the uniform distribution. Show that, for any $\delta > 0$, the following holds with

probability at least $1-\delta$ for all $q \in \{q\colon \|q-p\|_1 < 1\}$:

$$\underset{(x,y)\sim\mathcal{D}}{\mathbb{E}}[\ell(h(x),y)] \leq \sum_{t=1}^{T} q_t \ell(h_t(x_t), y_t) + 2\|q-p\|_1 + (\|q\|_2 + 2\|q-p\|_1)\left[2\sqrt{\log\log_2 \frac{2}{1-\|q-p\|_1}} + \sqrt{2\log\frac{2}{\delta}}\right].$$

Hint: consider the first inequality proven above for a fixed weight vector q^k and approximation error ϵ_k, for any $k \geq 0$. Show that the inequality can be extended to hold uniformly for all $k \geq 0$ if you choose $\epsilon_k = \epsilon + \sqrt{2\log(k+1)}$.

Solution: Consider two sequences $(\epsilon_k)_{k\geq 0}$ and $(q^k)_{k\geq 0}$. By the first part, for any fixed $k \geq 0$, we have

$$\mathbb{P}\left[\underset{(x,y)\sim\mathcal{D}}{\mathbb{E}}[\ell(h(x),y)] > \sum_{t=1}^{T} q_t^k \ell(h_t(x_t), y_t) + \|q^k\|_2\sqrt{2}\epsilon_k\right] \leq e^{-\epsilon_k^2}.$$

Choose $\epsilon_k = \epsilon + \sqrt{2\log(k+1)}$. Then, by the union bound, we can write:

$$\begin{aligned}&\mathbb{P}\left[\exists k \geq 1\colon \underset{(x,y)\sim\mathcal{D}}{\mathbb{E}}[\ell(h(x),y)] > \sum_{t=1}^{T} q_t^k \ell(h_t(x_t), y_t) + \|q^k\|_2\sqrt{2}\epsilon_k\right] \\ &\leq \sum_{k=0}^{+\infty} e^{-\epsilon_k^2} \leq \sum_{k=0}^{+\infty} e^{-\epsilon^2 - \log\left((k+1)^2\right)} = e^{-\epsilon^2}\sum_{k=1}^{+\infty}\frac{1}{k^2} = \frac{\pi^2}{6}e^{-\epsilon^2} \leq 2e^{-\epsilon^2}.\end{aligned} \tag{4}$$

We can choose q^k such that $\left\|q^k - p\right\|_1 = 1 - \frac{1}{2^k}$. Then, for any $q \in \{q\colon \|q - p\|_1 < 1\}$, there exists $k \geq 0$ such that $\left\|q^k - p\right\|_1 \leq \|q-p\|_1 < \left\|q^{k+1} - p\right\|_1$ and thus such that

$$\sqrt{2\log(k+1)} = \sqrt{2\log\log_2 \frac{1}{1-\left\|q^{k+1}-p\right\|_1}} = \sqrt{2\log\log_2 \frac{2}{1-\left\|q^{k}-p\right\|_1}} \leq \sqrt{2\log\log_2 \frac{2}{1-\left\|q-p\right\|_1}}$$

Furthermore, for that k, the following inequalities hold:

$$\begin{aligned}\left\|q^k\right\|_2 &\leq \|q\|_2 + \left\|q^k - q\right\|_2 \\ &\leq \|q\|_2 + \left\|q^k - q\right\|_1 \\ &\leq \|q\|_2 + \left\|q^k - p\right\|_1 + \|q-p\|_1 \\ &\leq \|q\|_2 + \|q-p\|_1 + \|q-p\|_1 \\ &\leq \|q\|_2 + 2\|q-p\|_1.\end{aligned}$$

and

$$\begin{aligned}\sum_{t=1}^T q_t^k \ell(h_t(x_t), y_t) &\leq \sum_{t=1}^T q_t \ell(h_t(x_t), y_t) + \left\|q^k - q\right\|_1 \\ &\leq \sum_{t=1}^T q_t \ell(h_t(x_t), y_t) + 2\|q-p\|_1.\end{aligned}$$

Plugging in these inequalities in (4) concludes the proof. $\square$

4 Coarse correlated equilibrium

Consider a finite normal form game with $p < +\infty$ players and finite action sets $\mathcal{A}_k$, $k \in [1, p]$. Show that if each player plays an external regret minimization strategy that has regret at most ϵ, then the empirical average of the players (product) distributions: $\mathsf{p} = \frac{1}{T}\sum_{t=1}^T \mathsf{p}^t$, where $\mathsf{p}^t = \prod_{k=1}^p \mathsf{p}_k^t$, is an ϵ-approximate coarse correlated equilibrium, that is, for all $k \in [1, p]$, for all $a_k \in \mathcal{A}_k$ and all $a'_k \in \mathcal{A}_k$,

$$\mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}}\left[u_k(a'_k, a_{-k})\right] \leq \mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}}[u_k(a_k, a_{-k})] + \epsilon.$$

Solution: By definition, for each player $k \in [1, p]$, the following holds:

$$\begin{aligned}\mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}}[u_k(a_k, a_{-k})] &= \frac{1}{T}\sum_{t=1}^T \mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}^t}[u_k(a_k, a_{-k})] \\ \mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}}\left[u_k(a'_k, a_{-k})\right] &= \frac{1}{T}\sum_{t=1}^T \mathop{\mathbb{E}}_{\mathsf{a}\sim\mathsf{p}^t}\left[u_k(a'_k, a_{-k})\right].\end{aligned}$$

Since each player has regret at most ϵ, for all $k \in [1, p]$,

$$\sup_{a'_k \in \mathcal{A}_k} \mathop{\mathbb{E}}_{\mathsf{a} \sim \mathsf{p}^t} \left[u_k(a'_k, a_{-k}) \right] \leq \mathop{\mathbb{E}}_{\mathsf{a} \sim \mathsf{p}^t} [u_k(a_k, a_{-k})] + \epsilon$$

This implies the ϵ-approximate coarse correlated equilibrium condition. $\square$