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 $\mathbb{E}[e^{tX}] \leq \cosh(t)$.

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 $R_T \leq 4\sqrt{\mathcal{L}_T^{\min} \log N}$.

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 $1/T$.

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 $q \in \{q\colon \|q - p\|_1 < 1\}$:

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

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 $\epsilon_k = \epsilon + \sqrt{2\log(k+1)}$.

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 $a'_k \in \mathcal{A}_k$,

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