Mehryar Mohri
Advanced Machine Learning 2017
Courant Institute of Mathematical Sciences
Homework assignment 1
February 27th, 2017
Due: March 20, 2017

A. Exponentially Weighted algorithm

The regret bound of the EW algorithm given in class does not match that of the Halving algorithm in the case where the loss of the best expert in hindsight is zero. In this problem, we will give a more favorable bound for such cases. We will adopt the same assumptions and use the same notation as for the regret bound theorem given in class for EW.

1. Fix $\eta > 0$. Show that for any $X \in [0, 1]$, $e^{-\eta X} \leq X(e^{-\eta} - 1) + 1$ (*hint*: use convexity of the exponential function). Use that to show that for a random variable X taking values in $[0, 1]$,
$$\log \mathrm{E}[e^{-\eta X}] \leq (e^{-\eta} - 1)\,\mathrm{E}[X]. \tag{1}$$

2. Prove that the cumulative loss of EW can be bounded as follows:
$$\sum_{t=1}^{T} L(\widehat{y}_t, y_t) \leq \frac{\eta L_T^*}{1 - e^{-\eta}} + \frac{\log N}{1 - e^{-\eta}}, \tag{2}$$
where L_T^* is the loss of the best expert in hindsight (*hint*: use inequality (1) instead of Hoeffding's inequality in the proof given in class). Compare this result with the Halving bound when $L_T^* = 0$ and η large.

3. Bonus question: Prove the inequality $\frac{\eta}{1-e^{-\eta}} \leq 1+\eta$. Use this to derive an upper bound and choose η to minimize that bound.

C. Correlated equilibria

Consider the following version of the Rock-Paper-Scissors where players are both penalized if they play the same action.

	R	P	S
R	(-1, -1)	(0, 1)	(1, 0)
P	(1, 0)	(-1, -1)	(0, 1)
S	(0, 1)	(1, 0)	(-1, -1)

1. Show that this game admits a unique mixed Nash equilibrium with non-zero probability for all actions. What is the expected payoff for the players?

2. Show that the game admits a correlated equilibrium with expected payoff $\frac{1}{2}$.

D. Mirror Descent

In class we presented a general guarantee for Mirror Descent. We will adopt the assumptions of that theorem as well as the notation. Additionally, assume that the functions f_t are β-strongly convex with respect to Φ, that is

$$f_t(\mathbf{w}') \geq f_t(\mathbf{w}) + \delta f_t(\mathbf{w}) \cdot (\mathbf{w}' - \mathbf{w}) + \frac{\beta}{2} B(\mathbf{w}' \parallel \mathbf{w}),$$

for all $\mathbf{w}', \mathbf{w} \in \mathbb{R}^N$, with $\beta > 0$.

1. Suppose we use Mirror Descent with a time-varying learning rate $\eta_{t+1} = \frac{2}{\beta t}$. Prove a logarithmic bound on the regret of Mirror Descent (*hint*: use proof for strongly convex losses given for PSGD). Your bound should be explicit and you should carefully justify all steps.

2. Show that your bound coincides with the one presented for PSGD in class in the case of strongly-convex losses.

B. Time-varying parameter

In this problem, we consider the use of the EW algorithm with no prior knowledge of the horizon T and using, instead of a fixed parameter η, a time-varying parameter $\eta_t > 0$, with $\eta_t \leq \eta_{t-1}$ for all $t \in [1, T]$. We define $w_{0,i} = 1$ for all $i \in [1, N]$. At iteration $t \geq 1$, prediction is made using the weights $w_{t-1,i}$ via

$$\widehat{y}_t = \frac{\sum_{i=1}^N w_{t-1,i}\, y_{t,i}}{\sum_{i=1}^N w_{t-1,i}}. \tag{3}$$

The weight of expert i is then updated as follows:

$$w_{t,i} = e^{-\eta_t L_{t,i}} \quad \text{with} \quad L_{t,i} = \sum_{s=1}^{t} L(\widehat{y}_{s,i}\,, y_s). \tag{4}$$

For any t, we also define $L_{t,*} = \min_{i \in [1,N]} L_{t,i}$ and $w_{t,*} = e^{-\eta_t L_{t,*}}$. We define W_t as the sum of the weights at time t: $W_t = \sum_{i=1}^N w_{t,i}$. Similarly, we define $w'_{t,i} = e^{-\eta_{t-1} L_{t,i}}$, $W'_t = \sum_{i=1}^N w'_{t,i}$, and $w'_{t,*} = e^{-\eta_{t-1} L_{t,*}}$.

For any $t \in [0, T]$, define the potential

$$\Phi_t = \frac{1}{\eta_t} \log \frac{W_t}{w_{t,*}}. \tag{5}$$

1. Show that the following equality holds for all $t \in [1, T]$:

$$\begin{aligned}\Phi_t - \Phi_{t-1} = & \underbrace{\frac{1}{\eta_t}\left[\log \frac{W_t}{w_{t,*}} - \log \frac{W'_t}{w'_{t,*}}\right]}_{A} + \underbrace{\left[\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}}\right] \log \frac{W'_t}{w'_{t,*}}}_{B} \\ & + \underbrace{\frac{1}{\eta_{t-1}}\left[\log \frac{W'_t}{w'_{t,*}} - \log \frac{W_{t-1}}{w_{t-1,*}}\right]}_{C}. \end{aligned} \tag{6}$$

2. Use the inequality:

$$\log \frac{\sum_{i=1}^N e^{-\eta_t [L_{t,i} - L_{t,*}]}}{\sum_{i=1}^N e^{-\eta_{t-1} [L_{t,i} - L_{t,*}]}} \leq \frac{\eta_{t-1} - \eta_t}{\eta_{t-1}} \log N, \tag{7}$$

   to show the following bound on A:

$$A \leq \left[\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}}\right] \log N. \tag{8}$$

   Bonus question: prove inequality (7) using Jensen's inequality and the convexity of the function $\Psi \colon \eta \mapsto \log \left[\sum_{i=1}^n e^{-\eta [L_{t,i} - L_{t,*}]}\right]$.

3. Show that B can be bounded as follows:

$$B \leq \left[\frac{1}{\eta_t} - \frac{1}{\eta_{t-1}}\right] \log N. \tag{9}$$

4. Use a technique similar to that of the proof of EW given in class to show that the third term can be bounded as follows:

$$C \leq L_{t,*} - L_{t-1,*} - L(\widehat{y}_t, y_t) + \frac{\eta_{t-1}}{8}. \tag{10}$$

5. Use the upper bounds on A, B, and C to prove the following upper bound on the regret of EW using a time-varying parameter:

$$R_T \leq \frac{1}{8}\sum_{t=1}^{T}\eta_{t-1} + \frac{2}{\eta_T}\log N - \frac{1}{\eta_0}\log N. \tag{11}$$

6. Assume now that for any $t \in [0, T]$, we choose $\eta_t = \sqrt{\frac{\alpha \log N}{t+1}}$, where $\alpha > 0$ is a parameter we will select. Show that the following upper bound holds:

$$R_T \leq \frac{1}{4}\sqrt{\alpha T \log N} + 2\sqrt{\frac{(T+1)\log N}{\alpha}} - \sqrt{\frac{\log N}{\alpha}}.$$

7. Prove the following regret bound using $\alpha = 8$:

$$R_T \leq \sqrt{2T\log N} + \sqrt{\frac{\log N}{8}}.$$

8. Bonus question: Derive a regret bound in terms of $L_{T,*}$ by choosing $\eta_t = \sqrt{\frac{\alpha \log N}{L_{t,*}+1}}$.