Mehryar Mohri
Advanced Machine Learning 2018
Courant Institute of Mathematical Sciences
Homework assignment 1
February 20, 2018
Due: March 06, 2018

A. Exponentially Weighted algorithm

We consider the Exponentially Weighted algorithm and adopt the notation and setup discussed in class. Let N_L be the number of experts with cumulative loss at most $L > 0$ at time T: $N_L = |\{i \in [N] \colon \sum_{t=1}^{T} L(\widehat{y}_{t,i}, y_t) \leq L\}|$.

1. Show that the following inequality holds for the cumulative loss of the algorithm, for any $\eta > 0$:

$$\sum_{t=1}^{T} L(\widehat{y}_t, y_t) \leq L + \frac{1}{\eta} \log \frac{N}{N_L} + \frac{\eta T}{8}.$$

Solution: Using the proof and notation presented in class, observe that the lower bound can be written as follows:

$$\begin{aligned}
\Phi_T - \Phi_0 = \log \sum_{i=1}^{N} e^{-\eta L_{T,i}} - \log N &\geq \log \sum_{L_{T,i} \leq L}^{N} e^{-\eta L_{T,i}} - \log N \\
&\geq \log N_L e^{-\eta L} - \log N \\
&= -\log \frac{N}{N_L} - \eta L.
\end{aligned}$$

Comparing this lower bound with the upper bound given in class yields

$$\begin{aligned}
&-\log \frac{N}{N_L} - \eta L \leq -\eta \sum_{t=1}^{T} L(\widehat{y}_t, y_t) + \frac{\eta^2 T}{8} \\
&\Rightarrow \sum_{t=1}^{T} L(\widehat{y}_t, y_t) \leq L + \frac{1}{\eta} \log \frac{N}{N_L} + \frac{\eta T}{8}.
\end{aligned}$$

2. Suppose L is very close to $\min_{i=1}^{N} L_{T,i}$ and that N_L is very large, say $N_L = N/2$. What does the bound show?

Solution: When a very large number of experts admit a cumulative loss close to the best in hinsight, the bound shows that the regret at time T of the algorithm is upper bounded by $\frac{1}{\eta}\log 2 + \frac{\eta T}{8}$, which is a significantly more favorable guarantee. FOr the best choice of η, the regret is bounded by the following: $\sqrt{2T\log(2)}$.

B. Games

1. Find all pure and mixed Nash equilibria of the following game.

	L	R
U	(0, 0)	(6, 2)
D	(2, 6)	(5, 5)

Solution: It is easy to check that both (D, L) and (U, R) are pure Nash equilibria. The sum of payoffs is 8 .

Let $c = (p, 1-p)$ be column player mixed strategy to play $\{U, D\}$, and let $r = (q, 1-q)$ be row player's mixed strategy to play $\{L, R\}$, where $0 < p < 1$ and $0 < q < 1$. If (c, r) is a mixed Nash equilibrium, then the expected payoff for column player to play L and R should be the same, otherwise the column player has incentive to remove probability from suboptimal action. Therefore,

$$6 * (1-q) = 2 * q + 5 * (1-q),$$

thus $q = \frac{1}{3}$. By the symmetry of the game, $p = \frac{1}{3}$. We have the unique mixed Nash equilibrium

$$c = \left(\frac{1}{3}, \frac{2}{3}\right), r = \left(\frac{1}{3}, \frac{2}{3}\right).$$

The sum of payoffs is $\frac{1}{9} * 0 + \frac{2}{9} * (6+2) * 2 + \frac{4}{9} * (5+5) = 8$.

2. Can you find a correlated equilibrium for which the sum of the players' payoffs is more favorable than that of any Nash equilibrium?

Solution: Consider the following correlated strategy:

	L	R
U	0	1/3
D	1/3	1/3

We show that it is a correlated equilibrium. Consider the payoffs of row player:

$$\begin{aligned}
&\mathsf{p}(U,L)u(U,L) + \mathsf{p}(U,R)u(U,R) = 2,\\
&\mathsf{p}(U,L)u(D,L) + \mathsf{p}(U,R)u(D,R) = \frac{5}{3} < 2,\\
&\mathsf{p}(D,L)u(D,L) + \mathsf{p}(D,R)u(D,R) = \frac{7}{3},\\
&\mathsf{p}(D,L)u(U,L) + \mathsf{p}(D,R)u(U,R) = 2 < \frac{7}{3}.
\end{aligned}$$

Therefore the row player has no incentive to deviate from the action recommended. By the symmetry of the game, the same result holds for the column player. Therefore this is a correlated equilibrium.

The sum of expected payoffs is $\frac{1}{3} * 8 * 2 + \frac{1}{3} * 10 = \frac{26}{3}$, which is more favorable than that of any Nash equilibrium.

3. Same two questions for the following game.

	L	C	R
U	(1, 1)	(2, 4)	(4, 2)
M	(4, 2)	(1, 1)	(2, 4)
D	(2, 4)	(4, 2)	(1, 1)

Solution: There is no pure Nash equilibrium. Let $c = (p_1, p_2, 1 - p_1 - p_2)$ be column player's mixed strategy to play $\{L, C, R\}$, and let $r = (q_1, q_2, 1 - q_1 - q_2)$ be row player's mixed strategy to play $\{U, M, D\}$. We first assume that all actions have positive probabilities. If (c, r) is a mixed Nash equilibrium, then column player should have same expected payoff for each action, which implies that

$$\begin{aligned}
&q_1 + 2q_2 + 4(1 - q_1 - q_2)\\
&= 4q_1 + q_2 + 2(1 - q_1 - q_2)\\
&= 2q_1 + 4q_2 + (1 - q_1 - q_2).
\end{aligned}$$

These linear equations admit a unique solution $q_1 = q_2 = \frac{1}{3}$. By symmetry, $p_1 = p_2 = \frac{1}{3}$. Thus,

$$c = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), r = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$$

is a Nash equilibrium, its sum of expected payoffs is $\frac{14}{3}$.

Next we show that any Nash equilibrium must have positive probabilities on all actions. Assume there exists a Nash equilibrium where column player put 0 probability on R. By symmetry, row player must put 0 probability on D. However, since $2q_1 + 4q_2 = 2(q_1 + 2q_2)$, column player must have twice expected payoff of playing R then that of playing L, therefore he has incentive to move probabilities from L to R, which contradicts the assumption of equilibrium. By the same argument, a Nash equilibrium must put positive probabilities on all actions. Thus $\big((1/3, 1/3, 1/3), (1/3, 1/3, 1/3)\big)$ is the unique Nash equilibrium.

It is easy to verify that the following is a correlated equilibrium.

	L	C	R
U	0	1/6	1/6
M	1/6	0	1/6
D	1/6	1/6	0

Its sum of expected payoffs is 6, which is greater than that of Nash equilibrium.

C. Alternative proof of the theorem of Nash

In class, we gave a full proof of the theorem of Nash. Here, we will give an alternative proof using Bregman divergences. We will adopt the notation and terminology introduced in class.

For any $k \in [p]$, let F_k be a strictly convex and differentiable function defined over an open and convex set C_k containing the simplex $\Delta(\mathcal{A}_k)$. We will denote by B_k the Bregman divergence associated to F_k.

1. For any $\mathsf{p} \in \Omega = \times_{k=1}^{p} \Delta(\mathcal{A}_k)$ and $k \in [p]$, we define the function $\mathsf{q} \mapsto \Psi_k(\mathsf{q}, \mathsf{p})$ over $\Delta(\mathcal{A}_k)$ by

$$\Psi_k(\mathsf{q}, \mathsf{p}) = -u_k(\mathsf{q}, \mathsf{p}_{-k}) + \mathsf{B}_k(\mathsf{q} \,\|\, \mathsf{p}_k).$$

Prove that, for any $\mathsf{p} \in \Omega$, $\min_{q \in \Delta(\mathcal{A}_k)} \Psi_k(\mathsf{q}, \mathsf{p})$ is attained.

Solution: Since F_k is differentiable, $\mathsf{q} \mapsto \mathsf{B}_k(\mathsf{q} \,\|\, \mathsf{p}_k)$ is also differentiable and therefore continuous. $\mathsf{q} \mapsto u_k(\mathsf{q}, \mathsf{p}_{-k})$ is linear thus continuous. Thus, $\mathsf{q} \mapsto \Psi_k(\mathsf{q}, \mathsf{p})$ is a continuous function as a sum of continuous functions and attains its minimum over the compact set $\Delta(\mathcal{A}_k)$.

2. Prove that the minimizer of $\Psi_k(\mathsf{q}, \mathsf{p})$ over $\Delta(\mathcal{A}_k)$ is unique (*hint*: show that $\mathsf{q} \mapsto \Psi_k(\mathsf{q}, \mathsf{p})$ is strictly convex).

Solution: Since F_k is strictly convex, $\mathsf{q} \mapsto \mathsf{B}_k(\mathsf{q} \,\|\, \mathsf{p}_k)$ is also strictly convex a sum of F_k and affine function. Thus, $\mathsf{q} \mapsto \Psi_k(\mathsf{q}, \mathsf{p})$ is also strictly convex a sum of $\mathsf{q} \mapsto \mathsf{B}_k(\mathsf{q} \,\|\, \mathsf{p}_k)$ and a linear function.

Now, suppose q and $\mathsf{q}' \neq \mathsf{q}$ are both minimizers. Then, for any $\alpha \in (0, 1)$, $\alpha\mathsf{q} + (1-\alpha)\mathsf{q}'$ is in $\alpha\mathsf{q} + (1-\alpha)\mathsf{q}'$ and

$$\begin{aligned}\Psi_k(\alpha\mathsf{q} + (1-\alpha)\mathsf{q}', \mathsf{p}) &< \alpha\Psi_k(\mathsf{q}, \mathsf{p}) + (1-\alpha)\Psi_k(\mathsf{q}', \mathsf{p})\\ &= \alpha\Psi_k(\mathsf{q}, \mathsf{p}) + (1-\alpha)\Psi_k(\mathsf{q}, \mathsf{p})\\ &= \Psi_k(\mathsf{q}, \mathsf{p}),\end{aligned}$$

which contradicts the minimality of $\Psi_k(\mathsf{q}, \mathsf{p})$. Thus, the minimizer is unique.

3. Let $f\colon \Omega \to \Omega$ be the function defined for any $\mathsf{p} \in \Omega$ by $f(\mathsf{p}) = (\mathsf{q}_1, \ldots, \mathsf{q}_p)$, with $\mathsf{q}_k = \operatorname{argmin}_{\mathsf{q} \in \Delta(\mathcal{A}_k)} \Psi_k(\mathsf{q}, \mathsf{p})$. Assume that, for any $\mathsf{p} \in \Omega$, $\operatorname{argmin}_{\mathsf{q} \in \Delta(\mathcal{A}_k)} \Psi_k(\mathsf{q}, \mathsf{p})$ is a continuous function of p.

   Show that f is well defined and show that f admits a fixed-point.

   Solution: f is well defined by the existence and uniqueness of q_k established in previous questions. The continuity of f is a direct consequence of the assumption. Thus, the existence of a fixed-point follows Brouwer's theorem.

4. Let $\mathsf{p} \in \Omega$ be a fixed-point of f. Show that for any $k \in [p]$ and $q \in \Delta(\mathcal{A}_k)$, the following inequality holds: $\Psi_k(\mathsf{q}, \mathsf{p}) \geq \Psi_k(\mathsf{p}_k, \mathsf{p})$ (*hint*: prove that the right-derivative of the function J defined over $[0, 1]$ by $J(\alpha) = \Psi_k(\alpha\mathsf{q} + (1-\alpha)\mathsf{p}_k, \mathsf{p})$ is non-negative). Prove the theorem of Nash.

   Solution: Observe that $J(0) = \Psi_k(\mathsf{p}_k, \mathsf{p}) = -u_k(\mathsf{p}_k, \mathsf{p}_{-k})$. For any $\alpha \in (0, 1]$, by definition of p as a fixed-point of f, and thus the minimizing property of p_k, we have

$$\frac{J(\alpha) - J(0)}{\alpha} = \frac{\Psi_k(\alpha\mathsf{q} + (1-\alpha)\mathsf{p}_k, \mathsf{p}) - \Psi_k(\mathsf{p}_k, \mathsf{p})}{\alpha} \geq 0.$$

   This implies the inequality $J'(0) \geq 0$. Now,

$$\begin{aligned}&[\mathsf{B}(\alpha\mathsf{q} + (1-\alpha)\mathsf{p}_k \,\|\, \mathsf{p}_k)]'_{\alpha=0}\\ &= [F(\alpha\mathsf{q} + (1-\alpha)\mathsf{p}_k) - F(\mathsf{p}_k) - \langle \nabla F(\mathsf{p}_k), \alpha(\mathsf{q} - \mathsf{p}_k)\rangle]'_{\alpha=0}\\ &= \langle \nabla F(\mathsf{p}_k), \mathsf{q} - \mathsf{p}_k\rangle - \langle \nabla F(\mathsf{p}_k), \mathsf{q} - \mathsf{p}_k\rangle\\ &= 0.\end{aligned}$$

$-u_k(\alpha \mathsf{q} + (1-\alpha)\mathsf{p}_k, \mathsf{p}_{-k})$ is a linear function of α, thus, we can write

$$J'(0) = -u_k(\mathsf{q} - \mathsf{p}_k, \mathsf{p}_{-k}) = -u_k(\mathsf{q}, \mathsf{p}_{-k}) + u_k(\mathsf{p}_k, \mathsf{p}_{-k}) \geq 0,$$

which proves the theorem of Nash.

5. Show that the function F_k defining the Bregman divergence $\mathsf{B}_k(\mathsf{q} \,\|\, \mathsf{p}_k) = \frac{1}{2}\|\mathsf{q} - \mathsf{p}_k\|_2^2$ satisfies the assumptions.

   Solution: The function $F_k(\mathsf{q}) = \frac{1}{2}\|\mathsf{q}\|_2^2$ is strictly convex on $\mathbb{R}^{n_k}$ and differentiable.

6. Let n_k be the cardinality of $\mathcal{A}_k = \{a_1, \ldots, a_{n_k}\}$ and let $\mathsf{v} \in \mathbb{R}^{n_k}$ be the vector whose jth coordinate is $u_k(a_j, \mathsf{p}_{-k})$, $j \in [n_k]$. Prove that

$$\operatorname*{argmin}_{\mathsf{q} \in \Delta(\mathcal{A}_k)} \Psi_k(\mathsf{q}, \mathsf{p}) = \operatorname*{argmin}_{\mathsf{q} \in \mathcal{A}_k} \frac{1}{2}\|\mathsf{q} - (\mathsf{p}_k + \mathsf{v})\|_2^2.$$

   Solution: Observe that $u(\mathsf{q}, \mathsf{p}_k) = \sum_{j=1}^{n_k} \mathsf{q}(a)\, u_k(a_j, \mathsf{p}_{-k}) = \mathsf{q} \cdot \mathsf{v}$. Thus,

$$\begin{aligned}
\Psi_k(\mathsf{q}, \mathsf{p}) &= -u(\mathsf{q}, \mathsf{p}_k) + \frac{1}{2}\|\mathsf{q} - \mathsf{p}_k\|_2^2 \\
&= -\mathsf{q} \cdot \mathsf{v} + \frac{1}{2}\|\mathsf{q}\|_2^2 - \mathsf{q} \cdot \mathsf{p}_k + \frac{1}{2}\|\mathsf{p}_k\|_2^2 \\
&= \frac{1}{2}\|\mathsf{q} - (\mathsf{p}_k + \mathsf{v})\|_2^2 + \frac{1}{2}\|\mathsf{p}_k\|_2^2 - \frac{1}{2}\|\mathsf{p}_k + \mathsf{v}\|_2^2.
\end{aligned}$$

   Since the last two terms do not depend on q, this proves the result.

7. [bonus question] Prove that $\mathsf{p} \mapsto \operatorname{argmin}_{\mathsf{q} \in \Delta(\mathcal{A}_k)} \Psi_k(\mathsf{q}, \mathsf{p})$ is continuous.