Algorithmic game theory – Tutorial 7*

1 Coarse correlated equilibria

For a normal-form game $G = (P, A, C)$ of n players, a probability distribution $p(a)$ on A is a *correlated equilibrium* in G if $\sum_{a_{-i} \in A_{-i}} C_i(a_i; a_{-i}) p(a_i; a_{-i}) \leq \sum_{a_{-i} \in A_{-i}} C_i(a'_i; a_{-i}) p(a_i; a_{-i})$ for every player $i \in P$ and all $a_i, a'_i \in A_i$. A probability distribution $p(a)$ on A is a *coarse correlated equilibrium* in G if $\sum_{a \in A} C_i(a) p(a) \leq \sum_{a \in A} C_i(a'_i; a_{-i}) p(a)$ for every player $i \in P$ and every $a'_i \in A_i$.

Exercise 1. *Show formally that every correlated equilibrium is a coarse correlated equilibrium.*

Exercise 2. *Compute all coarse correlated equilibria in the Prisoner's dilemma game.*

	T	S
T	(2,2)	(0,3)
S	(3,0)	(1,1)

Table 1: The game from Exercise 2.

2 Regret minimization

There are N available actions $X = \{1, \ldots, N\}$ and at each time step t the online algorithm A selects a probability distribution $p^t = (p_1^t, \ldots, p_N^t)$ over X. After the distribution p^t is chosen at time step t, the adversary chooses a loss vector $\ell^t = (\ell_1^t, \ldots, \ell_N^t) \in [-1, 1]^N$, where the number ℓ_i^t is the loss of action i in time t. The algorithm A then experiences loss $\ell_A^t = \sum_{i=1}^N p_i^t \ell_i^t$. After T steps, the loss of action i is $L_i^T = \sum_{t=1}^T \ell_i^t$ and the loss of A is $L_A^T = \sum_{t=1}^T \ell_A^t$. We use L_{min}^T to denote $\min_{i \in X} L_i^T$. The *external regret* of A is $R_A^T = \max_{i \in X}\{L_A^T - L_i^T\} = L_A^T - L_{min}^T$.

Exercise 3. *Assume an online algorithm A chooses among two actions, 1 and 2, over $T = 3$ steps. The losses for each action are given as:*

Step t	**Loss ℓ_1^t of 1**	**Loss ℓ_2^t of 2**
1	0.3	0.4
2	0.7	0.2
3	0.6	0.5

Table 2: Losses for each action over three steps from Exercise 3.

The algorithm A chooses the actions 1, 1, 2 in steps 1, 2, 3, respectively, all with probability 1. Compute the cumulative loss L_A^T of A, the cumulative loss L_i^T of always playing action $i \in \{1, 2\}$, and the external regret R_A^T of A.

Exercise 4. *Is the external regret always nonnegative? (That is, for any sequence $(\ell^t)_{t=1}^T$ of loss vectors and any algorithm A.) Can you come up with some upper bound on the external regret that always holds?*

Exercise 5. *An algorithm is* deterministic *if, for every step t, there is action i with $p_i^t = 1$. Show that for every deterministic algorithm D and $T \in \mathbb{N}$, there is a sequence of loss vectors such that $L_D^T = T$ and $L_{min}^T \leq \lfloor T/N \rfloor$. That is, $L_D^T \geq N \cdot L_{min}^T + (T \bmod N)$.*

Exercise 6 (*). *Prove the following statements about lower bounds on the external regret.*

*Information about the course can be found at http://kam.mff.cuni.cz/~sychrovsky/

(a) *For positive integers N and $T < \lfloor \log_2 N \rfloor$, consider the following sequence of loss vectors. At time step 1, a random subset of $N/2$ actions gets loss of 0 and the rest gets loss of 1. At time step $t \geq 2$, a random subset of half of the actions that have received loss 0 so far gets loss 0 and all remaining actions get loss of 1. Show that for every online algorithm A, we have $\mathbb{E}[L_A^T] \geq T/2$ and yet $L_{min}^T = 0$.*

(b) *In the case of $N = 2$ actions, consider the following sequence of loss vectors. Let $e_1 = (1, 0)$ and $e_2 = (0, 1)$. At each time step $t \in \{1, \dots, T\}$, we choose $\ell^t = e_1$ with probability 1/2 and $\ell^t = e_2$ with probability 1/2. Show that, for every online algorithm A, we have $\mathbb{E}[L_A^T - L_{min}^T] \geq \Omega(\sqrt{T})$.*