Algorithmic Game Theory: Problem Set 3

Due online via GradeScope before the start of class on Tuesday, March 9

Aaron Roth

Remember you can work together on problem sets, but list everyone you worked with, and everyone turn in their own assignment. Ask questions on Piazza.

Problem 1) No Deterministic No Regret (10 points)

In class, we proved that the polynomial weights algorithm achieves "no regret" – that is, for arbitrary sequences of losses, it guarantees that the difference between the average loss achieved by the algorithm, and the average loss achieved by the best expert in hindsight is $o(1)$ – tending to zero as $T \to \infty$. Recall that the polynomial weights algorithm is randomized. Here we show that no deterministic algorithm can obtain the same guarantee.

1. (5 pts) Consider the algorithm "Follow the Leader" that always picks the expert that has the lowest cumulative loss so far – i.e. at day j it picks expert k such that:

$$k = \arg\min_i \sum_{t=1}^{j-1} \ell_i^t.$$

   (Suppose for concreteness that if there is a tie, the algorithm picks the min-loss expert with the smallest index). Show that Follow the Leader is not a no-regret algorithm. i.e. exhibit a sequence of losses such that for all T, the regret of the algorithm is $\Omega(1)$.

2. (5 pts) Prove that no *deterministic* experts algorithm can achieve $o(1)$ regret – i.e. that randomization is necessary to achieve a guarantee like that of the polynomial weights algorithm.

   Hint: Consider any fixed deterministic algorithm, and then place yourself in the role of an adversary who is trying to foil it. Can you, knowing which expert the algorithm is going to pick next, design a sequence of losses so that after T rounds, the best expert always has cumulative loss that is lower than the algorithm's loss by at least $T/2$?

Problem 2) Multiplayer Zero Sum Games (15 pts)

Consider a two person zero-sum game in which each player has k actions, defined by an $k \times k$ matrix U. We can think of the actions of both players as being the sets $A_1 = A_2 = X = \{1, 2, \ldots, k\}$. Now consider the n person zero sum separable graphical game defined on a cycle of length n in which each person is simultaneously playing this zero-sum game with both their left neighbor and their right neighbor. With their right neighbor, they are taking the role of player 1 (the minimization player), and with their left neighbor, they are taking the role of player 2 (the maximization player). Formally, if player i is playing strategy s_i, then his utility is:

$$u_i(s) = U(s_{i-1}, s_i) - U(s_i, s_{i+1})$$

where we understand the index i mod n. (i.e. $0 \equiv n$ and $n+1 \equiv 1$).

1. (5 pts) Show that for every $\epsilon > 0$ any such game must have a symmetric ϵ-approximate Nash equilibrium – i.e. an ϵ-approximate Nash equilibrium in which every player is playing the exact same mixed strategy. (Hint: Think about what would happen if all of these players repeatedly play the game using the polynomial weights algorithm).

2. (5 pts) Show that in any symmetric Nash equilibrium, the expected utility of each player is 0.

3. (5 pts) Suppose instead of a cycle, the players are on a path (i.e. players 1 and n only participate in a single zero sum game, with players 2 and $n-1$ respectively). Is it still necessarily true that there is a Nash equilibrium in which every player has expected utility 0?

Problem 3) Almost Zero Sum Games (15 pts)

Consider a two player game G such that the two players have utility functions of the following form: for $a_1 \in A_1, a_2 \in A_2$:

$$u_1(a_1, a_2) = f(a_1, a_2) + g(a_1) \quad u_2(a_1, a_2) = -f(a_1, a_2) + h(a_2)$$

where f, g, h can be arbitrary functions. Observe that G need not be zero-sum. Define an alternative game $\hat{G}$ such that the two players have modified utility functions:

$$\hat{u}_1(a_1, a_2) = f(a_1, a_2) + g(a_1) - h(a_2) \quad \hat{u}_2(a_1, a_2) = -f(a_1, a_2) + h(a_2) - g(a_1)$$

Observe that $\hat{G}$ *is* zero-sum.

1. (5 pts) Show that (p, q) is a Nash equilibrium for $\hat{G}$ if and only if it is a Nash equilibrium for G. In other words, G and $\hat{G}$ are *strategically equivalent*.

2. (5 pts). Consider any sequence of action profiles $(a^1, \ldots, a^T)$. Show that if this sequence of action profiles has regret ϵ with respect to the player utility functions from G, then it also has regret ϵ with respect to the player utility functions from $\hat{G}$.

3. (5 pts) Conclude that if two players play using the polynomial weights algorithm in G, the empirical average of their play converges to Nash equilibrium (even though G is not 0 sum)

Problem 4) Symmetric Games (15 pts)

Consider a two person game in which each player has the same strategy set: $A_1 = A_2 = X$. Say that the game is *symmetric* if for all $a_1, a_2 \in X$:

$$u_1(a_1, a_2) = u_2(a_2, a_1)$$

1. (10 pts) Show that if there are only two actions (i.e. $|X| = 2$), any symmetric game must have a pure strategy Nash equilibrium.

2. (5 pts) Consider the 3 action case (i.e. $|X| = 3$). Either show that there must always exist a pure strategy Nash equilibrium, or give an example of a 3 action symmetric game that does not have any pure strategy Nash equilibrium.