Algorithmic Game Theory: Problem Set 3

Due online via GradeScope before the start of class on Tuesday, March 12

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 in $[0, 1]$ such that for all T, the average regret of the algorithm is $\Omega(1)$.

2. (5 pts) Prove that no *deterministic* experts algorithm can achieve $o(1)$ regret against arbitrary losses in $[0, 1]$ – 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) Regret to a Continuum of Experts (15 pts)

Polynomial weights assumes there are a finite number of experts n. However, there are settings where the number of "experts" may be infinite (e.g., prices in an interval). We modify the polynomial weights setting so that experts lie on a continuum, where expert $x \in [0, 1]$. For each $t \in \{1, \ldots, T\}$, $\ell_t : [0, 1] \to [0, 1]$ is an (adversarially chosen) function that maps an expert x to its loss in round t. We require that each ℓ_t is

L-Lipschitz continuous, which means that:

$$|\ell_t(x) - \ell_t(y)| \leq L \cdot |x - y| \qquad (\forall x, y \in [0,1], \text{ some } L \in \mathbb{R}_{>0})$$

1. (10 pts) Describe how this setting can be reduced to a finite number of experts by discretizing $[0, 1]$ into finitely many representative points at distance α from one another, and running polynomial weights on these representatives. What additional regret term, relative to the best expert in $[0, 1]$, would this incur when used in polynomial weights? How should we set the discretization parameter α? (Asymtptotic approximation is ok)

2. (5 pts) Show that no algorithm can get $o(1)$ regret to a continuum of experts if the adversary is not restricted to choosing Lipschitz continuous losses. (Hint: Give a concrete strategy for the Adversary that will force any randomized algorithm to obtain linear regret to the best expert in hindsight)

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) Correlated Equilibria (15 pts)

Recall the traffic light game:

	STOP	GO
STOP	(0,0)	(0,1)
GO	(1,0)	(-100,-100)

In class, we showed that this game has a correlated equilibrium (CE) of (GO, STOP)$\sim \frac{1}{2}$, (STOP, GO)$\sim \frac{1}{2}$ which attains better social welfare than any mixed Nash equilibrium (NE). Furthermore, the CE has the same social welfare as the game's two pure NE while being more fair than either. We will now prove that the social welfare of CE can sometimes strictly exceed that of *any* NE, pure or mixed.

Consider modifying the traffic light game so that players prefer stopping when the opponent is also stopping (perhaps it allows for some pleasant conversation), and the collision cost is significantly ameliorated (due to improvements in air bag technology):

	STOP	GO
STOP	$(\frac{2}{3}, \frac{2}{3})$	(0,1)
GO	(1,0)	(-1,-1)

1. (5 pts) Find a welfare-maximizing NE of the new game (taking into account both mixed and pure strategies). You must prove that there is no NE with higher welfare.
2. (10 pts) Prove that the welfare-maximizing CE of this game has strictly higher social welfare than any NE.