Lectures 6-7

Includes material from Arora, Hazan, Kale: "Experts: The Multiplicative Weights Update Method: a Meta Algorithm and Applications"

These are slightly modified from last semester's notes.

1 Overview

We will discuss the first of two important applications of the experts framework: approximately solving zero sum games.

1.1 Zero sum games

A zero sum game is specified by an $m \times n$ matrix A. The matrix entry A_{ij} is the loss for the row player if the row player plays i and the column player plays j. For this lecture, we assume that the losses belong to $[0, 1]$. The row player can follow a mixed strategy by choosing a row i according to a probability distribution on the rows. The column player can similarly follow a mixed strategy y. The expected loss for the row player is $x^t Ay$ if the players follow mixed strategies (x, y).

If the row player plays first and chooses strategy x, and the column player gets to choose an optimal response, then the loss for the row player is:

$$C(x) = \max_y x^t Ay$$

If the column player plays first and chooses strategy y, and the row player gets to choose an optimal response, then the loss for the row player is:

$$R(y) = \min_x x^t Ay$$

Von Neumann's minimax theorem states that $\max_y R(y) = \min_x C(x)$. The common value is called the value of the game and is denoted by V.

ϵ-optimal strategies: If x^* and y^* are best responses to each other, then $R(y^*) = C(x^*) = V$: the strategies are optimal. The column player's best response to a non-optimal strategy x has value more than V. We therefore have the inequality:

$$R(y) \leq V \leq C(x)$$

A pair of strategies is called ϵ-optimal if $C(x) - R(y) \leq \epsilon$.

2 Experts algorithm for zero sum games

Finding a pair of optimal strategies for a zero sum game can be reduced to solving a linear program, and the converse is also true. Finding ϵ-optimal strategies is a non trivial problem as it is like solving linear programs approximately.

We will assume that we can solve the simpler problem of finding the best response to a strategy played by the row player. If the payoff matrix is given explicitly, the best response is found by computing the payoff for every column and choosing the maximum.

$$\overline{C}(x) = \text{argmax}_{j \in [n]} (x^t A)_j$$

Even if the matrix A is implicitly specified and has an exponential number of columns, we will see that for some cases the optimal response can be computed in polynomial time.

2.1 Experts recap

Recall that the experts framework consists of n experts who each make predictions every day, and each expert incurs a loss that is revealed at the end of the day.

We analyzed the algorithm where all experts have weight 1 initially, the algorithm chooses expert i with probability proportional to the weight w_i, and weights are updated as $w_i(t+1) = w_i(t)(1-\epsilon)^{\ell_i(t)}$. The expected loss L for this algorithm is close to of the loss L^* of the best expert in retrospect,

$$L \leq (1+\epsilon)L^* + \frac{\ln n}{\epsilon} \tag{1}$$

2.2 Algorithm

We will use the experts algorithm to find an approximate equilbrium to a zero-sum game. The m pure strategies of the row player will be the experts. The experts algorithm picks experts according to a probability distribution; for the game setting choosing an expert probabilistically is equivalent to playing a mixed strategy.

An ϵ-approximate solution to a zero sum game can be found by iterating the following steps for $T = \frac{\log n}{\epsilon^2}$ rounds:

1. In round t, the row player plays the mixed strategy x_t specified by the experts algorithm.
2. The column player plays $j = \overline{C}(x_t)$, the optimal response to the row player's mixed strategy.
3. The loss of expert i is A_{ij} (the loss of the i-th pure row strategy against the column player's move).

The algorithm trains the row player by playing against a column player who always plays the best response. By following the experts algorithm, the row player ensures that the average over many rounds is close to an optimal strategy,

Theorem

The average strategies $x_{\text{avg}} = \frac{1}{T}\sum_t x(t)$ *and* $y_{\text{avg}} = \frac{1}{T}\sum_t y(t)$ *are a* 2ϵ*-optimal pair.*

Proof: For a zero sum game the column player's gain is the row player's loss, so the row player's total loss L is also the column player's total gain over T rounds.

The column player's average gain L/T would be $C(x_{\text{avg}})$ if the column player played $\overline{C}(x_{\text{avg}})$ in each round.

The column player chooses the best strategy in each round: in particular for round t the column player's gain for the chosen strategy $y(t)$ is at least the gain for $\overline{C}(x_{\text{avg}})$. Therefore L/T is at least $C(x_{\text{avg}})$, justifying the order of points in the following picture:

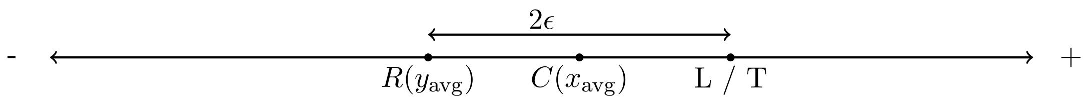

The best expert in retrospect is the best response to y_{avg}, so $R(y_{\text{avg}})$ is the average loss for the best expert. The analysis of the expert's algorithm (**??**) shows that the total loss L is not much worse than the loss $L^* = R(y_{\text{avg}})T$ of the best expert in retrospect:

$$L \leq (1+\epsilon)R(y_{\text{avg}})T + \frac{\ln n}{\epsilon} \Rightarrow L/T \leq R(y_{\text{avg}}) + 2\epsilon \tag{2}$$

We used the assumption that the losses lie in $[0, 1]$ for the inequality $\epsilon R(y_{\text{avg}}) \leq \epsilon$ and the choice of T for $\frac{\ln n}{\epsilon T} \leq \epsilon$. From the picture it follows that $C(x)_{\text{avg}} - R(y_{\text{avg}}) \leq 2\epsilon$, hence the strategies (x, y) are 2ϵ-optimal. □