CSE 522: Algorithms and Uncertainty	Spring 2017

Lecture 6 — April 14, 2017

Lecturer: Nikhil R. Devanur

In this class, we will discuss applications of the experts problem.

1 Zero sum games

A zero sum game is a simultaneous move game between 2 players. Such a game is represented by a matrix $A \in \mathbb{R}^{m\times n}$. The strategies of the "row" (resp. "column") player are the rows (resp. columns) of A. If the row player plays strategy $i \in [m]$ and the column player plays $j \in [n]$ then the outcome is A_{ij}.[1] Interpret this as that the row player pays A_{ij} amount of money to the column player, therefore the row player tries to minimize A_{ij} while the column player tries to maximize it.

The players may randomize their strategies; in this case we will let x and y denote the row and the column strategies respectively, with $x(i)$ and $y(j)$ denoting the probability of playing $i \in [m]$ and $j \in [n]$ resp. Let the simplex in $\mathbb{R}^d$ be denoted by

$$\Delta^d := \{x \in \mathbb{R}^d : \sum_i x(i) = 1, \forall i, x(i) \geq 0\}.$$

Then $x \in \Delta^m$ and $y \in \Delta^n$. A deterministic strategy that picks a row i (resp. column j) is represented by the unit vector $\mathbf{e}_i$ (resp. $\mathbf{e}_j$). The expected outcome of the game when the strategies are x and y is

$$x^\top Ay.$$

(We think of a vector in n dimensions as an $n \times 1$ matrix, and vice versa.) Suppose one of the players "moves first", which means that she announces her randomized strategy. The other player can then pick her strategy having known her opponents randomized strategy; we say that this player "moves second".

Exercise 1. *Show that for the player who moves second, randomized strategies are no better than deterministic strategies, i.e.,*

$$\forall x \in \Delta^m, \max_{y\in\Delta^n} x^\top Ay = \max_{j\in[n]} x^\top A\mathbf{e}_j.$$

$$\forall y \in \Delta^n, \min_{x\in\Delta^m} x^\top Ay = \min_{i\in[m]} \mathbf{e}_i^\top Ay.$$

For any given randomized strategy for the opponent, we call the strategy that optimizes the payoff as the *best response* to that strategy (i.e., the arg min/arg max in the above). Notation:

$$\lambda_{\min} := \min_{x\in\Delta^m} \max_{y\in\Delta^n} x^\top Ay.$$

$$\lambda_{\max} := \max_{y\in\Delta^n} \min_{x\in\Delta^m} x^\top Ay.$$

[1] The notation $[n]$ is a shorthand for the set $\{1, 2, \ldots, n\}$.

$\lambda_{\min}$ (resp. $\lambda_{\max}$) is the outcome when the row (resp. column) player moves first, and both players pick their strategies optimally. Moving second is always advantageous.This means that

$$\lambda_{\min} \geq \lambda_{\max}. \tag{1}$$

Exercise 2. *Prove Inequality (1).*

The main result regarding zero sum games is that moving first is not a disadvantage.

Theorem 1 (vonNeumann's Minmax Theorem)**.**

$$\lambda_{\min} = \lambda_{\max} = \lambda^*.$$

We call λ^* *as the* value *of the game.*

We will prove this theorem using the regret bound for the experts problem (Theorem 4 in Lecture 5).

1.1 Reduction from zero sum games to experts

Given an instance of a zero sum game (i.e., the payoff matrix A), we define an instance of the experts problem. The set of experts is the set of rows of A, i.e., the set $[m]$. Recall that in round t, the experts algorithm plays expert i with probability $w_t(i)/W_t$; this can be thought of as a randomized strategy for the row player. Denote the best response of the column player to this strategy by

$$j_t := \arg\max_{j \in [n]} w_t^\top A \mathbf{e}_j. \tag{2}$$

Define the losses in round t as the j_tth column of A, i.e.,

$$\ell_t(i) := A_{ij_t}.$$

Run the experts algorithm for T rounds. This completes the description of the instance of the experts algorithm.

Lemma 1. *In every round, the expected loss of the algorithm is at least as large as the min-max value. For all* $t \in [T]$,

$$\ell_t(\text{ALG}) \geq \lambda_{\min}.$$

Proof. Recall that $\lambda_{\min}$ is the value of the game when the row player moves first, and plays optimally. For any round t, we have that $\ell_t(\text{ALG})$ is the value of the game when the row player plays some strategy (given by w_t) and the column player best responds. This can therefore only be larger than $\lambda_{\min}$. □

Lemma 2. *The average loss of the best fixed expert/strategy on hindsight is no larger than the max-min value.*

$$\lambda_{\max} \geq \min_{i \in [m]} \sum_{t=1}^{T} \frac{1}{T} \ell_t(i).$$

Proof. The RHS corresponds to the payoff of the game when the column player moves first, plays the strategy corresponding to the empirical distribution of the j_ts, and the row player best responds. (The empirical distribution is the one that picks j_t with probability $1/T$, for each $t \in T$.) This is no larger than the payoff for the optimal column player strategy when he moves first, which is $\lambda_{\max}$. □

Proof of theorem 2. From Lemmas 1, 2 and Theorem 1, we have that $\lambda_{\min}$ and $\lambda_{\max}$ are sandwiched between the two terms that define the regret of the experts problem.

$$\frac{1}{T}\ell_{1..T}(\text{ALG}) \geq \lambda_{\min} \geq \lambda_{\max} \geq \frac{1}{T}\min_{i\in[m]} \ell_{1..T}(i).$$

From Theorem 4 in Lecture 5 we know that the difference between the first and the last term in the above sequence goes to 0 as $T \to \infty$. Hence the terms in the middle have to be the same. □

We now show how to actually read off approximately optimal strategies for the row and the column player. Suppose that T is such that $\text{REGRET}/T \leq \epsilon$. Then we will find row and column player strategies x' and y' such that

$$\max_{y\in\Delta^n} x'^{\top} Ay \leq \lambda^* + \epsilon.$$

$$\min_{x\in\Delta^m} x^{\top} Ay' \geq \lambda^* - \epsilon.$$

Note that for the column player the empirical distribution over the j_ts is approximately optimal. (This follows from the proof of Theorem 2.) For the row player, let t^* be the time index with the smallest loss for ALG:

$$t^* := \arg\min_{t\in T} \ell_t(\text{ALG}).$$

Define $x' = w_{t^*}/W_{t^*}$ and observe that $\max_{y\in\Delta^n} x'^{\top} Ay = \ell_{t^*}(\text{ALG})$. Once again, from the proof of Theorem 2, it follows that this is at most $\lambda^* + \epsilon$.

We will use a generalization of Theorem 2, where the strategies could be arbitrary convex sets. Let $K_1 \subseteq \mathbb{R}^m$ and $K_2 \subseteq \mathbb{R}^n$ be convex sets.[2]

Theorem 2 (vonNeumann's Minmax Theorem). *For any two convex sets $K_1 \subseteq \mathbb{R}^m$ and $K_2 \subseteq \mathbb{R}^n$, and a matrix $A \in \mathbb{R}_+^{m\times n}$,*

$$\min_{x\in K_1}\max_{y\in K_2} x^{\top} Ay = \max_{y\in K_2}\min_{x\in K_1} x^{\top} Ay = \lambda^*.$$

2 Linear Programs

Linear Programs (LPs) are widely used in optimization. We will consider a special form of an LP here, for the sake of simplicity. There are n *variables* in the LP, which we will denote by y. (The individual variables are $y(i), \forall\ i \in [n]$.) The input is a matrix $A \in \mathbb{R}^{m\times n}$, a vector $b \in \mathbb{R}^m$, and a convex set K. The set K is supposed to capture the "easy" constraints. This is formalized by

[2] A set K is convex iff for any two points $x, y \in K$, the entire line joining x and y is contained in K.

saying that we can optimize linear functions over K, i.e., we can solve the following optimization problem, for any given vector $c \in \mathbb{R}^n$.

$$\max_{y \in K} c \cdot y. \tag{3}$$

The input A and b capture the hard constraints which are of the form

$$Ay \geq b.$$

The problem is to determine whether all the constraints can be simultaneously satisfied or not. The output of the problem is supposed to be

- YES, if there exists a solution that satisfies $y \in K$ and $Ay \geq b$.
- NO, otherwise.

We will make the following assumption, due to which our LP will not be the most general. This is mostly for the sake of simplicity.

Assumption 1. *All co-ordinates of b are strictly positive. With this assumption, we can normalize each of the constraints and assume without loss of generality that b is the all ones vector.*

2.1 Reduction from LP to zero sum games

We will reduce this to a zero sum game. Define the following instance of a zero sum game, given an instance of an LP as above. The row player strategy set is $K_1 = [m]$, and the column player strategy set is $K_2 = K$. The payoff matrix is A. Let λ^* be the value of this game.

Lemma 3. *The LP instance is feasible if and only if $\lambda^* \geq 1$.*

Proof. Suppose that the LP is feasible. Let $y^* \in K$ be one feasible solution. Suppose the column player's strategy is y^*. Then no matter what $i \in [m]$ the row player picks, the ith coordinate of Ay^* is at least 1, since y^* is feasible. Therefore the value of the game is at least 1.

Suppose that the LP is infeasible. Let $y^* \in K$ be the optimal column player strategy. Then by definition,

$$\min_{i \in [m]} \mathbf{e}_i^\top A y^* = \lambda^*.$$

Since the LP is infeasible, there must exist at least one coordinate of Ay^* which is less than 1. Therefore the min in the above equation must be less than 1. □

Since we know how to solve zero-sum games using experts, we can in turn solve the LP using experts too. For this, we will further assume the following boundedness property: $\forall\, x \in K, Ax \in [0, a]$, for some given $a \in \mathbb{R}_+$. Consider combining the LP to zero sum game reduction, and the zero sum game to experts reduction. The set of experts will then be the rows of A. In order to define the $\ell_t(\cdot)$s, recall that we need to be able to solve the optimization problem as specified in (2), which in this case translates to the following optimization problem.

$$y_t = \arg\max_{y \in K} w_t^\top A y = \arg\max_{y \in K} (A^\top w_t) \cdot y.$$

This is exactly the kind of problem we assumed was easy to solve, in (3).

Consider a run of the experts problem on this instance given by the reduction.

- If in any round $t \in [T]$, it turns out that $l_t(\text{ALG}) < 1$, then output NO. This is because this implies (from Lemma 1) that $\lambda^* < l_t(\text{ALG}) < 1$, and then from Lemma 3 we know that the LP is infeasible.
- If in all rounds $t \in [T]$, $l_t(\text{ALG}) \geq 1$, then output $y' = \frac{1}{T}\sum_{t\in[T]} y_t$. Then from the regret bound, it follows that
$$\min_{i\in[m]} \mathbf{e}_i^\top Ay' \geq \frac{1}{T}\ell_{1..T}(\text{ALG}) - \epsilon \geq 1 - \epsilon.$$
This implies that y' is approximately feasible, i.e., it violates each constraint by at most ϵ.

3 Boosting

Recall the PAC learning setting: there is a distribution $\mathcal{D}$ over examples of feature and label pairs (x, y). The goal is to accurately predict the label y given the feature x. We define a "weak learner" and a "strong learner" as algorithms whose errors are slightly better than a random guess, and almost perfect, respectively.

Definition 1. *A* weak *(resp.* strong*) learner is an algorithm that given any distribution $\mathcal{D}$ (explicit, or with sample access), outputs a hypothesis h_{ALG} such that for some $\gamma > 0$ (resp. $\epsilon > 0$),*

$$err_{\mathcal{D}}(h_{\text{ALG}}) \leq \frac{1}{2} - \gamma \quad (\textbf{\textit{resp.}} \ \leq \epsilon).$$

In particular, suppose that we have a given training set of examples,

$$\{(x_1, y_1), (x_2, y_2), \ldots, (x_m, y_m)\},$$

and the distribution given to the strong learner is the empirical distribution over these examples.

For the weak learner we will assume that h_{ALG} is in some given set of hypothesis $\mathcal{H}$, and for the strong learner,
$$h_{\text{ALG}} = \text{WT-MAJORITY}(\mathcal{H}, \cdot, w(\cdot)),$$
for some set of weights $w(\cdot)$.

3.1 Reduction from Boosting to zero sum games

We would like to obtain a strong learner given a weak learner Why do we even hope to achieve this? A reduction to zero sum games shows this is possible, and the reduction to experts gives an algorithm!

Consider a zero sum game where the row player (deterministic) strategies are the hypothesis in $\mathcal{H}$, and column player strategies are the training set examples. The payoff of the game, given a hypothesis h and an example (x, y) is
$$\mathbf{1}(h(x) \neq y).$$

For any distribution on the training set $\mathcal{D}$, and any hypothesis $h \in \mathcal{H}$, when the row player strategy is $\mathcal{D}$ and the column player strategy is h, the expected payoff of the game is precisely $\text{err}_{\mathcal{D}}(h)$. Let λ^* as usual denote the value of this game. The weak learning guarantee implies a bound on λ^*: since we know that for all distributions $\mathcal{D}$, there exists an $h \in \mathcal{H}$ such that $\text{err}_{\mathcal{D}}(h) \leq \frac{1}{2} - \gamma$, it implies that

$$\lambda^* = \max_{\mathcal{D}} \min_{h \in \mathcal{H}} \text{err}_{\mathcal{D}}(h) \leq \frac{1}{2} - \gamma.$$

This now implies that there exists a probability distribution $\alpha(\cdot)$ over the hypothesis set $\mathcal{H}$ such that

$$\max_{i \in [m]} \mathbb{P}_{h \sim \alpha}[h(x_i) \neq y_i] = \lambda^* \leq \frac{1}{2} - \gamma.$$

In other words (or rather, symbols)

$$\forall i \in [m], \sum_{h \in \mathcal{H}} \alpha(h) \mathbf{1}(h(x_i) \neq y_i) \leq \frac{1}{2} - \gamma \leq \frac{1}{2} + \gamma \leq \sum_{h \in \mathcal{H}} \alpha(h) \mathbf{1}(h(x_i) = y_i),$$

since $\alpha(\cdot)$ is a probability distribution and therefore $\sum_{h \in \mathcal{H}} \alpha(h) = 1$. This means that the weighted sum of hypothesis that predict the correct label is always higher than those that predict the incorrect label. Hence

$$\forall i \in [m], \text{WT-MAJORITY}(\mathcal{H}, x_i, \alpha(\cdot)) = y_i.$$

This says that there is a weighted majority that predicts correctly on all the examples in the training set. Since this is an optimal strategy for a zero sum games, we can hope to actually find an approximately optimal strategy for this game using the experts algorithm.

Recall that in order to run the experts algorithm we need to be able to solve the optimization problem in (2). Since the row player strategy set is a probability distribution $\alpha(\cdot)$ over the set $\mathcal{H}$, and the column player strategy set is $[m]$, this translates to finding

$$\max_{i \in [m]} \mathbb{P}_{h \sim \alpha}[h(x_i) \neq y_i].$$

This doesn't seem like a problem we know how to solve. Moreover, the hypothesis set $\mathcal{H}$ maybe too large for us to run the experts algorithm efficiently.

Exercise 3. *Show how to solve the boosting problem using the experts algorithm.*