Stat 155 Lecture 9 Notes

Daniel Raban

February 15, 2018

1 Gradient Ascent, Series Games, and Parallel Games

1.1 Gradient ascent

Here, we will describe a low regret (in the sense that $R_T/T \to 0$ as $T \to \infty$) learning algorithm for a two player zero-sum game. This will complete our proof of the von Neumann minimax theorem.

Fix $x_1 \in \Delta_m$. On round t, play x_t, observe y_t, and choose

$$x_{t+1} = \mathcal{P}_{\Delta_m}(x_t + \eta A y_t),$$

where η us a step size and $\mathcal{P}_{\Delta_m}$ is the projection onto Δ_m:

$$\mathcal{P}_{\Delta_m}(x) = \operatorname*{arg\,min}_{a \in \Delta_m} \|a - x\|_2^2.$$

Note that if $F(x) = x^\top A y_t$, $\nabla F(x) = A y_y$. This is a "gradient ascent" algorithm because $A y_t$ is the gradient of the payoff when the column player plays y_t.

Theorem 1.1. *Let $G = \max_{y \in \Delta_n} \|Ay\|$. Then the gradient ascent algorithm with $\eta = \sqrt{2/(G^2 T)}$ has regret*

$$R_T \le \sqrt{2G^2 T}.$$

Proof. Note that

$$\begin{aligned} R_t &= \max_{x \in \Delta_m} \sum_{t=1}^T x^\top A y_y - \sum_{t=1}^T x_t^\top A y_t \\ &= \max_{x \in \Delta_m} \sum_{t=1}^T (x - x_t)^\top A y_y. \end{aligned}$$

Fix a strategy x. How does $\|x - x_t\|$ evolve?

$$\|x - x_{t+1}\| = \|x - \mathcal{P}_{\Delta_m}(x_t + \eta A y_t)\|$$

The distance to the projection is at most the distance to the original point.

$$\leq \|x - x_t - \eta A y_t\|$$

Use the identity that $\|a+b\|^2 = \|a\|^2 + 2a \cdot b + \|b\|^2$.

$$= \|x - x_t\|^2 - 2\eta (x - x_t)^\top A y_t + \eta^2 \|A y_t\|.$$

So we get that

$$2\eta (x - x_t)^\top A y_t \leq \|x - x_t\|^2 - \|x - x_{t+1}\|^2 + \eta^2 \|A y_t\|^2.$$

We can use this inequality to get

$$\begin{aligned}
\sum_{t=1}^{T} (x - x_t)^\top A y_t &\leq \frac{1}{2\eta} \sum_{t=1}^{T} (\|x - x_t\|^2 - \|x - x_{t+1}\|^2) + \frac{\eta}{2} \sum_{t=1}^{T} \|A y_t\|^2 \\
&= \frac{1}{2\eta} (\|x - x_1\|^2 - \|x - x_{T+1}\|^2) + \frac{\eta}{2} \sum_{t=1}^{T} \|A y_t\|^2 \\
&\leq \frac{2}{\eta} + \frac{\eta T G^2}{2}.
\end{aligned}$$

Choosing $\eta = \sqrt{2/(G^2 T)}$ and taking the max over x on the left side gives the result. $\square$

1.2 Series and parallel games

1.2.1 Series games

Say we have two games, G_1 and G_2. How can we combine these into a single game?

Definition 1.1. A *series game* is a game in which every turn, both players first play G_1 then both play G_2.

If the players play x_1 and y_1 in G_1 and then x_2 and y_2 in G_2, the payoff is $x_1^\top A y_1 + x_2^\top A_2 y_2$. The two games decouple; Player 1 should play x_1^* and x_2^*, and Player 2 should play y_1^* and y_2^*. If G_1 has value V_1, and G_2 has value V_2, the series game has value $V_1 + V_2$.

1.2.2 Parallel games

Definition 1.2. A *parallel game* is a game in which both players simultaneously decide which game to play, and an action in that game. If they choose the same game, they get the payoff from that game. If they choose different games, the payoff is 0.

Player 1 can either play x_1 in G_1 or x_2 in G_2. Player 2 can either play y_1 in G_1 or y_2 in G_2. If they both play G_1, the payoff is $x_1^\top A_1 y_1$. If they both play G_2, the payoff is $x_2^\top s A_2 y_2$. Otherwise, the payoff is 0. So the matrix for the game can be expressed as a block matrix:

$$\begin{pmatrix} A_1 & 0 \\ 0 & A_2 \end{pmatrix}.$$

We can split the decisions into choosing a mixture of games and then, with in each game, choosing a strategy. Withing G_1, Player 1 only needs to consider payoffs in G_1; if Player II chooses G_2, the payoff is 0, so Player 1 is indifferent about actions in that case. Thus, the players should play optimal strategies within each game, and the only choice is which game to play. So we can reduce the payoff matrix to involve V_1 and V_2 only:

$$\begin{pmatrix} V_1 & 0 \\ 0 & V_2 \end{pmatrix}.$$

We can solve this to find that Player 1 should play G_1 with probability

$$\frac{V_2}{V_1 + V_2}$$

and that the value of the game is

$$V = \frac{1}{1/V_1 + 1/V_2}.$$

What if we are playing k games in parallel? The payoff matrix becomes

$$\begin{pmatrix} V_1 & 0 & \cdots & 0 \\ 0 & V_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & V_k \end{pmatrix}.$$

If any entries are 0, this is a saddle point. If all entries are nonzero, the matrix is invertible and we can solve it by taking the inverse, as before. We also get

$$V = \frac{1}{1/V_1 + \cdots + 1/V_k}.$$

1.2.3 Electric networks

The way values combine in these games is identical to the way resistances combine in electric networks. For resistors connected in series, the *equivalent resistance* is the sum of the resistances of the resistors. For resistors connected in parallel, the equivalent resistance is the reciprocal of the sum of the reciprocals of the resistances.