Stat 155 Lecture 4 Notes

Daniel Raban

January 30, 2018

1 Two Player Zero-Sum Games

1.1 Pick a Hand

Consider a game of "Pick a Hand" with two players and two candies. The Hider puts both hands behind their back and chooses to either

1. Put 1 candy in their left hand (L_1),
2. Put 2 candies in their right hand (R_2).

The second player, the Chooser, picks a hand and takes the candies in it. Both moves are made simultaneously. We can represent this by a matrix:

	L_1	R_2
L	1	0
R	0	2

What if the players play randomly?

$$P(\text{Chooser plays } L) = x_1, \quad P(\text{Chooser plays } R) = 1 - x_1,$$

$$P(\text{Hider plays } L_1) = y_1, \quad P(\text{Chooser plays } R_2) = 1 - y_1.$$

Say we are playing sequentially, with the Chooser going first. The expected gain when the Hider plays L_1 is $x_1 \cdot 1 + (1 - x_1) \cdot 0 = x_1$. The expected gain when the Hider plays R_2 is $x_1 \cdot 0 + (1 - x_1) \cdot 2 = x(1 - x_1)$. Given these probabilities, the Holder can pick y_1 to minimize the Chooser's overall expected gain. The Chooser knows this, so the chooser should pick an x_1 that maximizes their expected gain given that they know that the Holder will minimize their expected gain. In this case, the Chooser should pick $x_1 = 2/3$. What if the Hider plays first? The Hider should also pick $y_1 = 2/3$.

1.2 Zero-sum games

Definition 1.1. A two player *zero-sum game* is a game where Player 1 has m actions $1, 2, \ldots, m$, and Player 2 has n actions $1, 2, \ldots, n$. The game has an $m \times n$ *payoff matrix* $A \in \mathbb{R}^{m \times n}$, which represents the payoff to player 1.

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \cdots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix}$$

If Player 1 chooses i, and Player 2 chooses j, then the payoff to player 1 is $a_{i,j}$, and the payoff to Player 2 is $-a_{i,j}$.

Definition 1.2. A *mixed strategy* is a probability distribution over actions. It is a vector

$$x = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_m \end{pmatrix} \in \Delta_m := \left\{ x \in \mathbb{R}^m : x_i \geq 0, \sum_{i=1}^{m} x_i = 1 \right\}$$

for Player 1 and

$$y = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix} \in \Delta_n := \left\{ x \in \mathbb{R}^n : y_i \geq 0, \sum_{i=1}^{n} y_i = 1 \right\}$$

for Player 2.

Definition 1.3. A *pure strategy* is a mixed strategy where one entry is 1, and all the others are 0. This is a standard basis vector e_i.

The expected payoff to Player 1 when Player 1 plays mixed strategy $x \in \Delta_m$ and Player 2 plays mixed strategy $y \in \Delta_m$ is

$$\begin{aligned} E_{I \sim x} E_{J \sim y} a_{I,J} &= \sum_{i=1}^{m} \sum_{j=1}^{n} c_i a_{i,j} y_j \\ &= x^\top A y \\ &= (x_1, x_2, \ldots, x_m) \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \cdots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}. \end{aligned}$$

Definition 1.4. A *safety strategy* for Player 1 is an $x^* \in \Delta_m$ that satisfies

$$\min_{y \in \Delta_n} (x^*)^\top Ay = \max_{x \in \Delta_m} \min_{y \in \Delta_n} x^\top Ay.$$

A safety strategy for Player 2 is an $y^* \in \Delta_n$ that satisfies

$$\max_{x \in \Delta_m} x^\top Ay^* = \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^\top Ay.$$

A safety strategy is the best strategy that Player 1 can use if they reveal their probability distribution to Player 2 before Player 2 makes a mixed strategy. This mixed strategy maximizes the worst case expected gain for Player 1. Safety strategies are optimal.

1.3 Von-Neumann's minimax theorem

Theorem 1.1 (Von-Neumann's Minimax Theorem)**.** *For any two-person zero-sum game with payoff matrix $A \in \mathbb{R}^{m \times n}$,*

$$\min_{y \in \Delta_n} \max_{x \in \Delta_m} x^\top Ay = \max_{x \in \Delta_m} \min_{y \in \Delta_n} x^\top Ay.$$

We will prove this in a later lecture. The left hand side says that Player 1 plays x first, and then Player 2 responds with y; the right hand side says that Player 2 plays y first, and then Player 1 responds with x.

You might think that this is actually an inequality ($\geq$) instead of an equality; this means playing last is preferable. But the minimax theorem says that it doesn't matter whether you play first or second.

Definition 1.5. We call the optimal expected payoff the *value* of the game.

$$V = \min_{y \in \Delta_n} \max_{x \in \Delta_m} x^\top Ay = \max_{x \in \Delta_m} \min_{y \in \Delta_n} x^\top Ay.$$