

APPLICATIONS OF LP: ZERO-SUM GAMES - Part 2

Last video: max min, min max over pure strategies

This video: Mixed / randomized strategies.

Example:

Min

Rock Paper Scissors

Max	Rock	0	-1	+1
	Paper	+1	0	-1
	Scissors	-1	+1	0

M_{RPS}

Rock
Paper
Scissors

Payoff for Maximizer

$$\text{max-min-pure } (M_{RPS}) = -1$$

$$\text{min-max-pure } (M_{RPS}) = +1$$

Consider a different strategy for Max:

- Max plays each strategy with $\frac{1}{3}$ Probability

Mixed Strategy

$$\sigma : \frac{1}{3} \text{ Rock} + \frac{1}{3} \text{ Paper} + \frac{1}{3} \text{ Scissors}$$

If Min plays pure strategy **Rock**, "Expected Payoff" = $\frac{1}{3} \cdot 0 + \frac{1}{3} \cdot 1 + \frac{1}{3} \cdot (-1)$
Paper, = $\frac{1}{3} \cdot (-1) + \frac{1}{3} \cdot 0 + \frac{1}{3} \cdot 1$
Scissors, = $\frac{1}{3} \cdot 1 + \frac{1}{3} \cdot (-1) + \frac{1}{3} \cdot 0$

When Max plays mixed strategy σ and Min plays some pure strategy, the expected payoff = 0

- But Min can also play mixed strategies. Suppose Min plays:

$$\tau : \frac{1}{3} \text{ Rock} + \frac{1}{3} \text{ Paper} + \frac{1}{3} \text{ Scissors}$$

- What is the expected payoff when σ and τ are played?

	$\frac{1}{3}$ Rock	$\frac{1}{3}$ Paper	$\frac{1}{3}$ Scissors
$\frac{1}{3}$ Rock	0	-1	+1
$\frac{1}{3}$ Paper	+1	0	-1
$\frac{1}{3}$ Scissors	-1	+1	0

$$\frac{1}{9} [(0 \quad -1 \quad +1) + (+1 \quad 0 \quad -1) + (-1 \quad +1 \quad 0)] = 0$$

Row 1 Row 2 Row 3

Mixed Strategy: is a probability distribution over pure strategies.

Maximizer's pure strategies: $\{1, 2, 3, \dots, m\}$

Mixed strategy for Max: $x_1, x_2, \dots, x_m \in \mathbb{R}_{\geq 0}$

$$\text{s.t. } \sum_{i=1}^m x_i = 1$$

Minimizer's pure strategies: $\{1, 2, 3, \dots, n\}$

Mixed strategy for Min: $y_1, y_2, \dots, y_n \in \mathbb{R}_{\geq 0}$

$$\text{s.t. } \sum_{j=1}^n y_j = 1$$

Payoff: Given $\sigma := (x_1, x_2, \dots, x_m)$ and $\tau := (y_1, y_2, \dots, y_n)$

$$\text{Payoff}(\sigma, \tau) = \sum_{i \in \{1, \dots, m\}} \sum_{j \in \{1, \dots, n\}} x_i \cdot y_j \cdot m_{ij}$$

$$\hookrightarrow = x^T M y$$

We are interested in the following quantities:

$$\max \min (M) = \max_{\substack{\text{mixed strategy } \sigma \\ \text{of Max}}} \min_{\substack{\text{mixed strategy } \tau \\ \text{of Min}}} \text{Payoff}(\sigma, \tau)$$

$$\min \max (M) = \min_{\substack{\text{mixed strategy } \tau \\ \text{of Min}}} \max_{\substack{\text{mixed strategy } \sigma \\ \text{of Max}}} \text{Payoff}(\sigma, \tau)$$

Notice that there are infinitely many mixed strategies.

Goal of this video: Given M , how do we compute

$\max \min (M)$?

- This does not look like a linear optimization problem since there is a mix of objectives and the cost that needs to be optimized: $x^T M y$ has bilinear terms

- We will now see that computing $\max \min (M)$ can be reduced to an LP problem.

1. Suppose we **fix** a strategy $\langle x_1, x_2, \dots, x_m \rangle$ for Max

- For example: $\langle 1/3, 1/3, 1/3 \rangle$ in MRPS.

2. Given the strategy $\langle x_1, \dots, x_m \rangle$ we want to find:

$$\min_{\text{mixed strategies of Min}} x^T M y$$

3. This can be written as a linear program (as x_i 's are constants)

$$\begin{array}{ll} \text{minimize} & x^T M y \\ \text{subject to} & y_1 + y_2 + \dots + y_n = 1 \\ & y_1, y_2, \dots, y_n \geq 0 \end{array} \quad (*)$$

For example, with $\langle 1/3, 1/3, 1/3 \rangle$ in RPS:

$$\begin{array}{ll} \text{minimize} & -1/3 y_2 + 1/3 y_3 + 1/3 y_1 - 1/3 y_3 - 1/3 y_1 + 1/3 y_2 \\ \text{subject to} & y_1 + y_2 + y_3 = 1 \\ & y_1, \dots, y_3 \geq 0 \end{array}$$


4. The answer to above LP gives the **best response** of Min to Max's mixed strategy $\langle x_1, x_2, \dots, x_m \rangle$

Each $\langle x_1, x_2, \dots, x_m \rangle \rightarrow x_0$ (answer of above LP)

We want the maximum possible x_0 .

5. Let us first write the dual q (*):

$$\begin{aligned} & \text{minimize} && x^T M y \\ & \text{subject to} && y_1 + y_2 + \dots + y_n = 1 \\ & && y_1, y_2, \dots, y_n \geq 0 \end{aligned}$$



$$\begin{aligned} & \max c^T x \\ & Ax \leq b \\ & x \geq 0 \end{aligned} \quad \leftarrow \quad \begin{aligned} & \min b^T y \\ & A^T y = c \\ & y \geq 0 \end{aligned}$$

$$\text{maximize } x_0$$

$$\begin{aligned} \text{subject to: } & x_0 \leq M_{11} x_1 + M_{21} x_2 + \dots + M_{m1} x_m \\ & x_0 \leq M_{12} x_1 + M_{22} x_2 + \dots + M_{m2} x_m \\ & \vdots \\ & x_0 \leq M_{1n} x_1 + M_{2n} x_2 + \dots + M_{mn} x_m \end{aligned}$$

$$\begin{aligned} & \text{maximize} && x_0 \\ & \text{subject to} && \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix} x_0 \leq M^T x \end{aligned} \quad (**)$$

$n \text{ rows} \rightarrow$
 $n \times 1$

6. Primal is bounded (feasible region is contained in unit hypercube)

Hence: Optimum of primal = optimum of dual

- For a given strategy of Max $\langle x_1, x_2, \dots, x_m \rangle$, the payoff obtained when Min plays her best response is given by the optimum cost of LP (**).

7. We want to maximize the optimum of LP (***) w.r.t. all mixed strategies of Maximizer.

- Hence we consider $\langle x_1, x_2, \dots, x_m \rangle$ as variables and add the constraint $\sum_{i=1}^m x_i = 1$ to (***)

$$\begin{array}{ll}
 \text{maximize} & x_0 \\
 \text{subject to} & \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} x_0 \leq M^T x \\
 & x_1 + x_2 + \dots + x_m = 1 \\
 & x_1, x_2, \dots, x_m \geq 0
 \end{array}$$

LP for max min (M) [over mixed strategies]

Illustration on the RPs example:

		Rock	Paper	Scissors									
<p>maximize x_0</p> <p>subject to:</p> <p>$x_0 \leq 0 \cdot x_1 + 1 \cdot x_2 - 1 \cdot x_3$</p> <p>$x_0 \leq -1 \cdot x_1 + 0 \cdot x_2 + 1 \cdot x_3$</p> <p>$x_0 \leq +1 \cdot x_1 - 1 \cdot x_2 + 0 \cdot x_3$</p> <p>$x_1 + x_2 + x_3 = 1$</p> <p>$x_1, x_2, x_3 \geq 0$</p>	<p>Rock</p> <p>Paper</p> <p>Scissors</p>	<table border="1" style="border-collapse: collapse; width: 100%;"> <tr> <td style="padding: 2px 5px;">0</td> <td style="padding: 2px 5px;">-1</td> <td style="padding: 2px 5px;">+1</td> </tr> <tr> <td style="padding: 2px 5px;">+1</td> <td style="padding: 2px 5px;">0</td> <td style="padding: 2px 5px;">-1</td> </tr> <tr> <td style="padding: 2px 5px;">-1</td> <td style="padding: 2px 5px;">+1</td> <td style="padding: 2px 5px;">0</td> </tr> </table>	0	-1	+1	+1	0	-1	-1	+1	0		
0	-1	+1											
+1	0	-1											
-1	+1	0											

← LP to find max min.