LINEAR PROGRAMMING

[V. CH11]: MATRIX GAMES

Phillip Keldenich Ahmad Moradi

Department of Computer Science
Algorithms Department
TU Braunschweig

February 7, 2023

Matrix Games

A two person game like the famous kid's Rock-Paper-Scissors game.

Rules.
At the count of three declare one of: Rock / Paper / Scissors

Winner Selection.
Identical selection is a draw. Otherwise:

- Rock dulls Scissors
- Paper covers Rock
- Scissors cuts Paper

Payoffs are from row player to column player:

$$\begin{array}{cc} & \begin{array}{ccc} R & P & S \end{array} \\ \begin{array}{c} R \\ P \\ S \end{array} & \begin{pmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{pmatrix} \end{array}$$

which one has the edge in (a round) of this game?
total payoff sent/received = 3/3

Note:
Any deterministic strategy employed by either player can be defeated systematically by the other player.

How about this one?

$$\begin{array}{c} \\ R \\ P \\ S \end{array} \begin{array}{c} \begin{array}{ccc} R & P & S \end{array} \\ \begin{pmatrix} 0 & 1 & -2 \\ -3 & 0 & 4 \\ 5 & -6 & 0 \end{pmatrix} \end{array}$$

total payoff sent/received = 10/11

Note:
Any deterministic strategy employed by either player can be defeated systematically by the other player.

Considering total payoff, one might suspect that *row player* might have the edge.

But

- Is this correct?
- if so, how much can the row player *expect* to win *on average* in each round?

Note:
How the row player might make use of this expected value if he knows it but the column player does not?

Lets first define such a game in general:

Given: an $m \times n$ matrix $A = [a_{ij}]$ known to both players in advance.

- Row player selects a *strategy* $i \in \{1, \ldots, m\}$
- Column player selects a strategy $j \in \{1, \ldots, n\}$
- Row player pays column player a_{ij} dollars.

Note: The rows of A represent deterministic strategies for row player, while columns of A represent deterministic strategies for column player.

Deterministic strategies can be (and usually are) bad.

Lets consider that they play with *randomised strategies*.

- Suppose column player picks j with probability x_j.
- Suppose row player picks i with probability y_i.

vectors $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_m)$ are stochastic vectors i.e. they have non-negative components that sums up to 1.

$$\forall j : x_j \geq 0 \quad \text{and} \quad \sum_j x_j = 1$$

Similarly

$$\forall i : y_i \geq 0 \quad \text{and} \quad \sum_i y_i = 1$$

If row player uses random strategy y and column player uses x, then expected payoff from row player to column player is

$$\sum_i \sum_j y_i a_{ij} x_j = y^T A x$$

Recall: For outcome (i, j) the payoff is a_{ij}, and, assuming that the row and column players behave *independently*, the probability of this outcome is simply $y_i x_j$.

Suppose column player were to adopt strategy x.

Then, row player's best defense is to use strategy y that *minimises* $y^T Ax$:

$$\min_y y^T Ax$$

And so column player should choose that x which *maximises* these possibilities:

$$\max_x \min_y y^T Ax$$

For a moment let $x = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})^T$.

Write down the inner minimisation problem.
Why this is only needed to look at deterministic strategies?

The observation gives a general way to solve the inner problem.

Given a vector x, the inner minimisation problem is easy as

$$\min_y y^T Ax = \min_i e_i^T Ax$$

e_i is the vector of all zeros except for a 1 in the i-th position: the deterministic strategy i.

Note: this observation reduced a minimisation over a continuum to one over a finite set.

We have:

$$\max(\min_i e_i^T Ax)$$

$$\sum_j x_j = 1$$

$$x_j \geq 0, \quad j = 1, \ldots, n$$

Introduce a scalar variable $v =$ *the value of the inner minimisation*, we get

$$\begin{aligned} &\max v \\ &v \geq e_i^T Ax, \qquad i = 1, \ldots, m \\ &\sum_j x_j = 1, \\ &x_j \geq 0, \qquad\quad j = 1, \ldots, n \end{aligned}$$

Writing in matrix-vector notation:

$$\begin{aligned} \max v& \\ ve - Ax &\leq 0 \\ e^T x &= 1 \\ x &\geq 0 \end{aligned}$$

(e without a subscript denotes the vector of all ones).

Similarly, row player seeks y^* attaining:

$$\min_y \max_x y^T A x$$

which is equivalent to:

$$\begin{aligned} \min u \\ ue - A^T y \geq 0 \\ e^T y = 1 \\ y \geq 0 \end{aligned}$$

Note:
Row player's problem is dual to column player's

column player's problem

$$\max \quad \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix}$$

$$s.t. \quad \begin{bmatrix} -A & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} x \\ v \end{bmatrix} \begin{matrix} \leq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$x \geq 0$$

$$v \text{ free}$$

row player's problem

$$\min \quad \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix}$$

$$s.t. \quad \begin{bmatrix} -A^T & e \\ e^T & 0 \end{bmatrix} \begin{bmatrix} y \\ u \end{bmatrix} \begin{matrix} \geq \\ = \end{matrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

$$y \geq 0$$

$$u \text{ free}$$

Theorem

There exist stochastic vectors x^* and y^* for which

$$\max_x {y^*}^T Ax = \min_y y^T Ax^*$$

The common optimal value $v^* = u^*$ of the primal and dual linear programs is called the value of the game.

What is the value of the game for the following matrix game

$$\begin{pmatrix} 0 & 1 & -2 \\ -3 & 0 & 4 \\ 5 & -6 & 0 \end{pmatrix}$$

In the tutorial session this week, we will discuss how linear programming could act as a classifier, i.e. the *support vector machine*.