LECTURE 12: GAME THEORY

Today: A taste of Game Theory and its relation to LP problems.

1. Example 1: Rock, Paper, Scissors

Let's play our favorite game: Rock, Paper, Scissors!

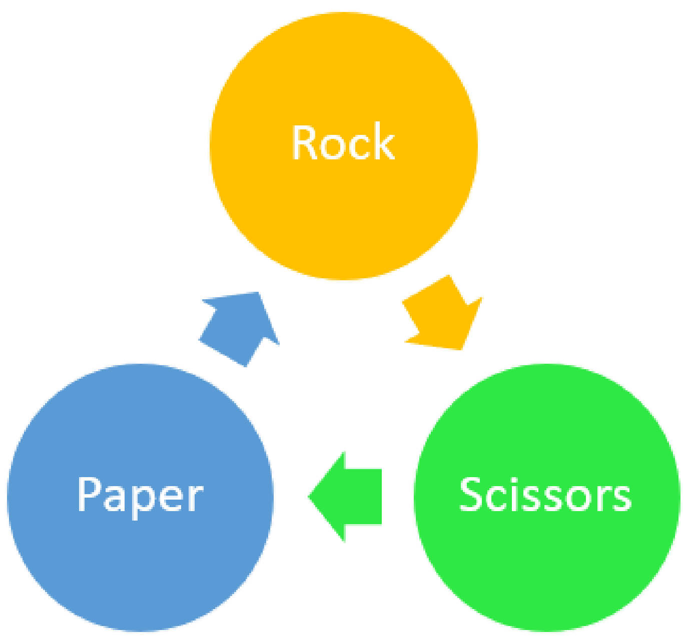

Two Players: Player 1 (P1) and Player 2 (P2)

Strategy: Each player picks a strategy: Rock, Paper, or Scissors.

Date: Tuesday, October 18, 2022.

Outcome: The winner gets \$5, the loser gets $-\$5$ and if it's a tie they both get \$0.

Zero-Sum Game: Here the gain for one player equals to the loss for the other player, this is called a **zero sum game**

Non-Example: "Not zero sum" would mean that P1 would win \$5 but P2 would only lose \$3

Because of this, it is enough to focus on Player 1 because if P1 gets \$ 5, we automatically know that P2 gets $-\$5$ and vice-versa.

Gains Matrix: Can represent this game with a matrix G, where the Rows correspond to P1 and the Columns correspond to P2. The $(i, j)-$th entry is the outcome for P1.

In this case, here is what the matrix G of gains looks like:

P1 \ P2	Rock	Paper	Scissor
Rock	0	-5	5
Paper	5	0	-5
Scissor	-5	5	0

Here is how to read this: If P1 chooses Paper (Row 2) and P2 chooses Rock (Column 1), this corresponds to the $(2, 1)$ entry, which is 5, so P1 gets \$5, and we automatically know that P2 gets $-\$5$ since it's a Zero-Sum Game.

2. Strategy

Goal: Figure out what strategy P1 and P2 should use.

Example: A stupid strategy to use is for P1 to always choose "Rock" because once P2 starts to notice the pattern, P2 will start choosing "Paper" and P1 will keep losing.

Mixed Strategy: The best thing is for P1 to choose a **mixed** strategy, where for example they will do Rock $\frac{1}{4}$th of the time, Paper $\frac{1}{2}$ of the time, and Scissors $\frac{1}{4}$th of the time.

More generally, assume P1 has a mixed strategy $x = (x_1, x_2, x_3)$ where $x_1, x_2, x_3 \geq 0$ and $x_1 + x_2 + x_3 = 1$. Here x_1 is the probability of P1 choosing Rock, x_2 the probability of choosing

Now P2 has itself a strategy, which we denote by $y = (y_1, y_2, y_3)$ with $y_i \geq 0$ and $y_1 + y_2 + y_3 = 1$

Each Round: P1 chooses Row i via the strategy vector $x = (x_1, x_2, x_3)$ and P2 chooses Row j via the vector $y = (y_1, y_2, y_3)$.

The probability of picking entry (i, j) is then $x_i y_j$ with gain G_{ij}.

Example: If $x_2 = \frac{1}{2}$ and $y_1 = \frac{1}{3}$ then the probability of the $(2, 1)-$th entry is $\frac{1}{2} \times \frac{1}{3} = \frac{1}{6}$ with gain 5

In practice we don't play one round, but several rounds, so it makes sense to talk about the average (or expected) pay-off, which just means summing up all the numbers above:

Definition: (Expected Payoff for P1)

$$\sum_{i,j} G_{ij} x_i y_j$$

Note: Compare this with Expected Value from Probability, which is $\sum$ values $\times$ probability

Note: P1 wants to **maximize** this expected payoff (maximize the gains) whereas P2 wants to **minimize** it (minimize the damage done)

Example 1:

Calculate the expected payoff in the case where $x = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$

You could do it by calculating the double-sum, but here it's easier to do it by cases, depending on P2's strategy.

Case 1: P2 chooses Rock, with probability y_1

Then the average payoff is (look at the first column of G)

$$\frac{1}{3} \times 0 + \frac{1}{3} \times 5 + \frac{1}{3} \times (-5) = 0$$

Case 2: P2 chooses Paper, same

Case 3: P2 chooses Scissors, same

Therefore, the expected payoff is

$$y_1 \times 0 + y_2 \times 0 + y_3 \times 0 = 0$$

Notice this is **independent** of P2's strategy! Similarly the average payoff of P2 is 0 regardless of P1's strategy.

Note: This trick works here because the x_i are all the same and because the columns of G sum to 0. For a proof, we have:

$$\sum_{i,j} G_{ij} \underbrace{x_i}_{\frac{1}{3}} y_j = \sum_{i,j} \frac{1}{3} G_{ij} y_j = \sum_j \sum_i \frac{1}{3} G_{ij} y_j = \sum_j y_j \underbrace{\left(\sum_i \frac{1}{3} G_{ij} \right)}_{=0} = 0$$

(In the last step, we used that the column sums are 0)

Let's now look at a case where the matrix is not symmetric.

3. Example 2: Pokémon Battle

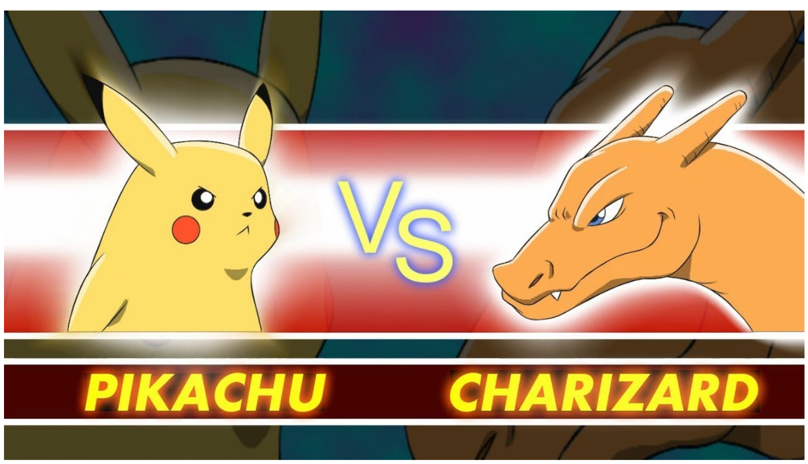

Two Players: Pikachu (P1) vs. Charizard (P2)

Strategy: Pikachu uses "Slam" or "Thunder" and Charizard uses "Fire Ball" or "Air Attack"

Gains Matrix:

Pikachu \ Charizard	Fire	Air
Slam	3	-1
Thunder	-2	1

Note: It's still zero-sum because if Pickachu gains 3 hit points, then Charizard loses 3 hit points.

Example 2:

Suppose that $x = \left(\frac{1}{2}, \frac{1}{2}\right)$ which strategy should P2 pick?

Case 1: If Charizard uses Fire Ball, then the average payoff is

$$\frac{1}{2} \times 3 + \frac{1}{2} \times (-2) = \frac{1}{2}$$

This means Charizard *loses* $\frac{1}{2}$

Case 2: If Charizard uses Air Attack, then the average payoff is

$$\frac{1}{2} \times (-1) + \frac{1}{2} \times 1 = 0$$

So it would make sense for Charizard to pick Air Attack, since it gives the smallest loss/damage

That said, we *don't* a priori know Pikachu's strategy, so we have to generalize:

Suppose Pikachu's strategy is $x = (x_1, x_2)$

Then Case 1 becomes $3x_1 - 2x_2$ and Case 2 becomes $-x_1 + x_2$ and therefore Charizard's strategy is to pick the *smaller* one of $3x_1 - 2x_2$ and $-x_1 + x_2$, so it makes sense to consider

$$\min \{3x_1 - 2x_2, -x_1 + x_2\}$$

(Before we had $\min\left\{\frac{1}{2}, 0\right\} = 0$ so we chose Air Attack)

4. Relation to LP

$$\text{Let } z = \min \{3x_1 - 2x_2, -x_1 + x_2\}$$

z is Charizard's best defense against Pickachu's strategy (x_1, x_2)

Note: This is *equivalent* to requiring

$$\begin{aligned} z &\leq 3x_1 - 2x_2 \\ z &\leq -x_1 + x_2 \end{aligned}$$

Because the minimum of two functions f and g is smaller than both f and g, and also because of those conditions will always be attained with $=$ by definition of a vertex.

But what does *Pikachu* want to do? It wants to maximize the damage done to Charizard, that is to pick x_1 and x_2 to maximize the damage z, and therefore Pikachu's goal is to solve the following

Pikachu's LP Problem:

$$\begin{aligned} &\max\ z \\ \text{subject to } & -3x_1 + 2x_2 + z \leq 0 \\ & x_1 - x_2 + z \leq 0 \\ & x_1 + x_2 = 1 \\ & x_1, x_2 \geq 0 \end{aligned}$$

Which is a LP problem!!!

Note: Notice the really cool change of point of view here! We started with Charizard's problem and then turned it into Pikachu's problem. It makes sense because Pikachu needs to think like Charizard.

5. Charizard's strategy

We have settled Pikachu's strategy, now what about Charizard's?

In this case, if Charizard's strategy is $y = (y_1, y_2)$ then

Case 1: Pikachu uses Slam, then the average payoff is $3y_1 - y_2$

Case 2: Pikachu uses Thunder, then the average payoff is $-2y_1 + y_2$

Note: Here we're looking at the *rows* instead of the columns

Therefore Pikachu's strategy is to pick the *bigger* one of the two:

$$w = \max\{3y_1 - y_2, -2y_1 + y_2\}$$

Which again is equivalent to

$$\begin{aligned} w &\geq 3y_1 - y_2 \\ w &\geq -2y_1 + y_2 \end{aligned}$$

And Charizard wants to minimize Pikachu's gains, that is pick y_1 and y_2 to minimize w. Hence Charizard's goal is to solve the following LP

$$\begin{aligned} &\min \ w \\ \text{subject to } & -3y_1 + y_2 + w \geq 0 \\ & 2y_1 - y_2 + w \geq 0 \\ & y_1 + y_2 = 1 \\ & y_1, y_2 \geq 0 \end{aligned}$$

Which is **PRECISELY** the dual problem!!!

In other words, the strategies of Pikachu and Charizard (P1 and P2) are dual to each other, **WOW!!!**

Note: It makes sense now why we consider A^T in the dual problem, because the columns of A^T correspond to the rows of A, and here for Pikachu's payoff we looked at the rows of A

Remark: The Strong Duality Theorem tells you that the values of both LP's are the same, that is $\max z = \min w$. Since z is a min (damage) and w is a max (gain), this becomes

$$\max_x \underbrace{\min_y \ \text{Payoff}}_{z} = \min_y \underbrace{\max_x \ \text{Payoff}}_{w}$$

Which is sometimes called the **minimax** theorem in Game theory

Note: Think of the min-max theorem as kind of a tug-of-war situation between the two players. Unless $\max \min > \min \max$ we create an imbalance that clearly makes one or the other the winner. If $\max \min = \min \max$, then we have an equilibrium situation.

6. MATLAB Implementation

Let's solve both the LP and the dual using MATLAB.

Remarks:

(1) For help with linprog, type "help linprog" in MATLAB

(2) The Lambda feature in linprog can actually solve for the dual problem directly!

LP Problem:

$$\begin{aligned} &\max\ z \\ \text{subject to}\ &-3x_1+2x_2+z\leq 0 \\ &x_1-x_2+z\leq 0 \\ &x_1+x_2=1 \\ &x_1,x_2\geq 0 \end{aligned}$$

To implement it into MATLAB, we need to rewrite this as

$$\begin{aligned} &\max\ 0x_1+0x_2-z=0 \\ \text{subject to}\ &-3x_1+2x_2+z\leq 0 \\ &x_1-x_2+z\leq 0 \\ &x_1+x_2\leq 1 \\ &-x_1-x_2\leq -1 \\ &-x_1\leq 0 \\ &-x_2\leq 0 \end{aligned}$$

The middle part follows since $x_1+x_2=1$ is equivalent to $x_1+x_2\leq 1$ and $-x_1-x_2\leq -1$

```
f = [0 0 -1];
A = [-3 2 1; 1 -1 1; 1 1 0; -1 -1 0; -1 0 0; 0 -1 0];
b = [0;0;1;-1;0;0];
[X,FVAL,EXITFLAG,OUTPUT,LAMBDA] = linprog(f,A,b);
```

If you type in X, then you get

```
X =
0.4286
0.5714
0.1429
```

Which tells you that the optimal strategy for Pikachu is $(x_1, x_2) \approx (0.4286, 0.5714)$ (the exact values are $\left(\frac{3}{7}, \frac{4}{7}\right)$) and the optimal z value is 0.1429 (exact value $z = \frac{1}{7}$)

You can even find the optimal strategy for the dual by typing in LAMBDA.ineqlin which gives you

```
ans =
0.2857
0.7143
0.1429
0
0
0
```

This tells you that the optimal strategy for Charizard is $(y_1, y_2) \approx (0.2857, 0.7143)$ which is $\left(\frac{2}{7}, \frac{5}{7}\right)$ and the optimal w value is 0.1429 that is $\frac{1}{7}$ (as expected)