MINIMAX LP FORMULATION

Luis Castellanos, Naisha Sarna, Andrej Jakovljevic

December 5 2025

1 Intro

Game theory is the study of mathematical models of strategic interactions. Moreover, it analyzes situations in which two or more decision-makers, called players, make choices that affect one another. A game, in this context, refers to when players contend with each other according to a list of rules.

Game theory provides a mathematical framework for studying strategies among multiple players. In such environments, each player's outcome depends not only on their own decisions but also on the decisions made by others. A strategy is a complete and predetermined plan describing how a player will act in every possible circumstance of the game. A strategy may be pure (always selecting the same action) or mixed (randomizing over available actions according to assigned probabilities).

When players select actions from a finite set, each combination of actions produces an outcome. A numerical value, called the payoff, is assigned to each outcome and reflects the reward or loss obtained by a player. Payoffs are commonly represented in a payoff matrix, where each entry quantifies the gain or loss resulting from the player's actions.

For this project, we focus on two-player zero-sum games. A zero-sum game is one in which the total payoff across players sum up to zero. Thus, any gain for one player corresponds exactly to a loss for the other. Because of this structure, we can describe the entire interaction using a single payoff matrix A. Each entry A_{ij} represents the payoff to Player1 (P1) when P1 selects action i while P2 selects action j. Since the game is zero-sum, P2's payoff is exactly $-A_{ij}$.

In many zero-sum settings, relying solely on pure strategies is not optimal. If a player repeats the same action, the opponent may exploit this predictability. This motivates the use of mixed strategies, defined as probability distributions over a player's available actions. By randomizing, a player avoids predictability and therefore avoids systematic exploitation.

The minimax principle formalizes optimal play in adversarial environments. Under minimax, each player selects a strategy that maximizes their own guaranteed payoff assuming the opponent behaves in the most hostile manner possible. For P1, this means maximizing the minimum payoff they can be forced to receive and for P2, it means minimizing the maximum payoff that P1 can secure. This principle forms the foundation of equilibrium analysis in zero-sum games and has significant influence in economics, optimization, artificial intelligence, and reinforcement learning.

2 A Motivating Example

To motivate the minimax framework, consider the classical zero-sum game Rock, Paper, and Scissors (RPS). In RPS, if P1 wins they get +1 and P2 gets -1, and vice versa and if they tie they both get 0, so every possible outcome adds up to zero. Therefore we conclude that RPS is a zero-sum game and we can proceed to define the payoff matrix for

P1 as:

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix},$$

where rows and columns correspond to Rock, Paper, and Scissors, respectively.

Note that if P1 plays Rock every time, then P2 can always respond with Paper and achieve a guaranteed payoff of 1. Hence, any fixed pure strategy is exploitable.

Let P1 choose a mixed strategy, we denote x_i as the probability of choosing i so in our example we'll have the following:

$$x = (x_R, x_P, x_S), \qquad x_R + x_P + x_S = 1, \quad x_i \geq 0.$$

Next, we will calculate the expected payoffs for P1 if P2 plays:

- Rock (column 1) is
$$(Ax)_1 = 0 \cdot x_R + 1 \cdot x_P + (-1) \cdot x_S = x_P - x_S,$$
- Paper (column 2) is
$$(Ax)_2 = (-1) \cdot x_R + 0 \cdot x_P + 1 \cdot x_S = -x_R + x_S,$$
- Scissors (column 3) is
$$(Ax)_3 = 1 \cdot x_R + (-1) \cdot x_P + 0 \cdot x_S = x_R - x_P.$$

If P1 wants to guarantee a payoff of at least v, these three values must each be at least v, thus we get the following constraints:

$$\begin{aligned} x_P - x_S &\geq v, \\ -x_R + x_S &\geq v, \\ x_R - x_P &\geq v, \end{aligned}$$

together with $x_R + x_P + x_S = 1$ and $x_R, x_P, x_S \geq 0$.

The minimax objective for P1 is

$$\max_{x_R, x_P, x_S} \min\{x_P - x_S,\ -x_R + x_S,\ x_R - x_P\},$$

Let us formulate the primal LP:

$$\begin{aligned} \max_{x_R, x_P, x_S, v} \quad & v \\ \text{subject to} \quad & x_P - x_S \geq v, \\ & -x_R + x_S \geq v, \\ & x_R - x_P \geq v, \\ & x_R + x_P + x_S = 1, \\ & x_R, x_P, x_S \geq 0. \end{aligned}$$

Note that at optimum all three payoff constraints are tight, meaning:

$$x_P - x_S = v, \quad -x_R + x_S = v, \quad x_R - x_P = v.$$

So we can solve this system of equalities and $x_R + x_P + x_S = 1$. We get the following solution:

$$x_R = x_P = x_S = \frac{1}{3}.$$

To find v, we use, $v = x_P - x_S$:

$$v = \frac{1}{3} - \frac{1}{3} = 0.$$

Hence the LP yields the optimal mixed strategy and game value:

$$x^\star = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \qquad v^\star = 0.$$

What these results we can conclude the following: (1) P1's optimal strategy is to play Rock, Paper and Scissors with equal probability so no option can be exploited, (2) the value of $v^\star = 0$ means that when both players play optimally, neither has an advantage and the expected payoff is zero.

The dual LP is:

$$\begin{aligned} \min_{y_R, y_P, y_S, w} \quad & w \\ \text{subject to} \quad & y_P - y_S \le w, \\ & -y_R + y_S \le w, \\ & y_R - y_P \le w, \\ & y_R + y_P + y_S = 1, \\ & y_R, y_P, y_S \ge 0. \end{aligned}$$

We can solve in the same way as in the primal, and verify that the dual LP agrees with the primal LP.

$$y^\star = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right), \qquad w^\star = 0.$$

A miracle happens the primal and dual have the same optimal values, in class we learned that this is no miracle this is strong duality and if we apply strong duality to zero-sum games we obtain the minimax theorem.

3 Generalization to Arbitrary Zero-Sum Games

Any finite two-player zero-sum game can be represented by an $m \times n$ payoff matrix A. P1 chooses $x \in \mathbb{R}^m$ and P2 chooses $y \in \mathbb{R}^n$,

$$x_i \ge 0, \ \sum_{i=1}^{m} x_i = 1, \qquad y_j \ge 0, \ \sum_{j=1}^{n} y_j = 1.$$

The expected payoff is $x^\top Ay$.

P1 wants to solve:

$$\max_x \min_j (Ax)_j,$$

and P2 wants:

$$\min_y \max_i (A^\top y)_i.$$

The Minimax Theorem guarantees that both expressions have equal value:

$$\max_x \min_y x^\top Ay = \min_y \max_x x^\top Ay,$$

and that optimal mixed strategies exist.

3.1 Primal LP Formulation

For a general payoff matrix $A \in \mathbb{R}^{m \times n}$, let $x \in \mathbb{R}^m$ be P1's mixed strategy and let v be the guaranteed payoff. The LP for P1 is

$$\begin{aligned} \max_{x,v} \quad & v \\ \text{subject to} \quad & (Ax)_j \ge v \quad \text{for all } j = 1, \dots, n, \\ & \sum_{i=1}^{m} x_i = 1, \\ & x_i \ge 0. \end{aligned}$$

$$(Ax)_j = \sum_{i=1}^{m} A_{ij} x_i,$$

3.2 Dual LP Formulation

We now look at the dual LP, which corresponds to P2's optimization problem. In general, the dual of the primal minimax LP for Player 1 is

$$\begin{aligned} \min_{y,w} \quad & w \\ \text{subject to} \quad & (A^\top y)_i \leq w \quad \text{for all } i = 1, \ldots, m, \\ & \sum_{j=1}^{n} y_j = 1, \\ & y_j \geq 0. \end{aligned}$$

4 The Colonel Blotto Game

The Colonel Blotto Game is a zero=sum game where two players (who are described as ”generals”) must allocate limited resources simultaneously across many battlefields. These resources can span from soldiers to money to even just units. The goal of this game is to allocate more resources to the same battlefield compared to your opponent; the total payoff is how many battlefields are won.

Players don't know how many resources their opponent is allocating to each specific battlefield, so the player must develop a strategy that can win against most of the opponents battlefields.

While this may seem very military focused, it has many real life examples. Some which include: political campaigns, security allocation, sport strategies, marketing, college admission, and more. Since players can place many resources into many battlefields, the number of allocations grows as the battlefields/resources grow, making finding the optimal solution almost impossible to solve by hand.

For a general game, let there be n battlefields indexed by $j = 1, 2, ..., n$. Two players; A and B, have their budgets, B_A and B_B respectively. Players A and B choose an allocation vector $\mathbf{a} = (a_1, ..., a_n), \mathbf{b} = (b_1, ..., b_n)$ with constraints:

$$a_j \geq 0, \sum_{j=1}^{n} a_j \leq B_A, b_j \geq 0, \sum_{j=1}^{n} b_j \leq B_B,$$

A the payoff is in the ”winner takes all” style, which gives player A's utility as:

$$U_A(\mathbf{a}, \mathbf{b}) = \sum_{j=1}^{n} v_j \cdot 1\{a_j > b_j\}$$

Where v_j is the value of battlefield j. The utility of B is typically the negative (since it is zero sum) and is $-U_A$.

4.1 Political Campaign

The specific example we plan to use is the political campaign problem, to be exact, a presidential campaign in the US. Since there are 51 states plus DC, these will be our ”battlefields”. The ”resources” will be the monetary budget allocated to each state. Each state also has a set about of electoral votes, making winning certain states more important than winning others.

We have written an algorithm to simulate an election. Of course, this is not entirely accurate as the allocations

	State	EV	A_alloc	B_alloc	Winner
0	CA	54.0	3.412728	1.407855	A
1	TX	40.0	1.110657	2.010129	B
2	FL	30.0	0.414348	2.786872	B
3	NY	28.0	0.603656	0.088955	A
4	PA	19.0	1.424167	2.632197	B
5	IL	19.0	0.786446	4.508921	B
6	OH	17.0	0.600838	3.694538	B
7	GA	16.0	0.947459	4.262212	B
8	NC	16.0	1.828709	0.994354	A
9	MI	15.0	0.190601	4.440243	B
10	NJ	14.0	2.028216	1.822077	A
11	VA	13.0	1.687203	1.926019	B
12	WA	12.0	2.033207	4.110250	B
13	AZ	11.0	3.720511	4.199744	B
14	IN	11.0	0.247520	0.574171	B

A total EV: 249.0
B total EV: 289.0

Figure 1: Example allocations and payoffs

per state are randomized and in real life there would be more strategic approaches based on different parameters, but this is a basic example. As there are 51 battlefields and many resources, it is almost impossible to do by hand.

`https://github.com/naibai101/MA-421-Colonel-Blotto-Algo.git`

In our algorithm our code shows the top 15 most important states by electoral votes, the allocations each player made to each of those states, and which player won each state. At the bottom it shows the total payoff which is calculated by adding all the electoral votes that each player won.

5 Resources

https://courses.cs.duke.edu/compsci570/fall19/lpandgames.pdf

https://vanderbei.princeton.edu/542/lectures/lec8.pdf

https://abel.math.harvard.edu/archive/20spring06/handouts/Lesson24-GameTheoryandLPhandout.pdf

https://people.csail.mit.edu/rivest/pubs/BBDHx18.pdf