Game Theory

Greg Plaxton
Theory in Programming Practice, Spring 2004
Department of Computer Science
University of Texas at Austin

Bimatrix Games

- We are given two real $m \times n$ matrices $A = (a_{ij})$, $B = (b_{ij})$, where $1 \leq i \leq m$ and $1 \leq j \leq n$
- There are two players, a *row player* and a *column player*
- The row player chooses a row i, and the column player chooses a column j
  - Each player's choice is made without knowledge of the other player's choice
- The payoff to the row player is a_{ij}, and the payoff to the column player is b_{ij}
- What is a good strategy for playing such a game?
  - This is a classic problem in game theory

Zero-Sum Games

- In this lecture we will focus primarily on the special case of a bimatrix game in which $B = -A$, i.e., the total payoff to the row and column players is always zero
  - These are called zero-sum games
  - Since B can be determined from A, we can consider the input to be the single matrix A

Example: Rock-Paper-Scissors

- Rock beats scissors, scissors beats paper, paper beats rock
- The winner gets a payoff of 1, and the loser gets a payoff of -1
- If both players play the same thing (e.g., rock), the payoff to each player is 0
- What is an optimal strategy for playing this game?

Mixed Strategy

- A *mixed strategy* for the column player is a probability distribution over the columns
  - Rather than deterministically picking a particular column, the column player fixes a probability distribution over the columns and then selects at random from this distribution
  - If the distribution assigns probability 1 to a particular column, it is a *pure strategy*
- Similarly, a *mixed strategy* for the row player is a probability distribution over the rows
- What is a good mixed strategy for the rock-paper-scissors game?
  - Is there a sense in which this strategy is optimal?

Zero-Sum Games: Can Assume $A \geq 0$

- Note that a_{ij} represents the payoff from the column player to the row player in the case where the row player plays row i and the column player plays column j
- We can assume without loss of generality that $A \geq 0$, i.e., the column player always pays a nonnegative amount
  - To see this, note that the structure of the problem is unchanged if we add some real value Δ to every a_{ij}
  - By choosing Δ sufficiently large, we can ensure that all of the a_{ij}'s are nonnegative
- We make this assumption throughout the remainder of the lecture

Expected Payoff

- Let A be the $m \times n$ payoff matrix for a zero-sum game
- Let $x = \langle x_1, \ldots, x_n \rangle$ denote the mixed strategy of the column player
  - The column player plays column j with probability x_j
  - Note that $\sum_{1 \le j \le n} x_j = 1$ and all of the x_j's are nonnegative
- Similarly, let $y = \langle y_1, \ldots, y_m \rangle$ denote the mixed strategy of the row player
- The expected payoff from the column player to the row player is

$$P(x,y) = \sum_{1 \le i \le m} \sum_{1 \le j \le n} x_j \cdot y_i \cdot a_{ij}$$

A Notion of Optimality for the Column Player

- Let x be an arbitrary mixed strategy for the column player
- Let $f(x)$ denote a mixed strategy for the row player that maximizes $P(x, f(x))$
- We say that x is *optimal* if it minimizes $P(x, f(x))$
  - Such an optimal mixed strategy is called a minimax strategy
- How can we efficiently compute a minimax strategy for the column player?
- Symmetrically, how can we efficiently compute a maximin strategy for the row player?

Computation of a Minimax Strategy

- Observation: For every mixed strategy x of the column player, there is a pure strategy y for the row player maximizing $P(x, y)$
  - Suppose the strategy y maximizing $P(x, y)$ is mixed and that $y_i > 0$
  - Then the pure strategy y' that always plays row i satisfies $P(x, y') = P(x, y)$
- Accordingly, we can formulate the optimization problem for the column player as follows
  - Determine a mixed strategy x and a (minimax) payoff α such that α is minimized and the inequality

    $$\sum_{1 \leq j \leq n} x_j \cdot a_{ij} \leq \alpha$$

    holds for all rows i
  - Is this a linear program?

Feasibility of the Minimax LP

- Note that the minimax LP is feasible and has a finite optimal value for the objective function α
  - Any mixed strategy x, coupled with a sufficiently large choice for α, yields a feasible solution
  - The sum of the a_{ij}'s is a trivial upper bound on the optimal value of the objective function

The Maximin LP

- Similarly, we can formulate an LP to determine an optimal mixed strategy for the row player

- Determine a mixed strategy y and a (maximin) payoff β such that β is maximized and the inequality $\left(\sum_{1\leq i\leq m} y_i \cdot a_{ij}\right) - \beta \geq 0$ holds for all columns j
  - The variables are the y_i's and β
  - The requirement that y is a mixed strategy is enforced by the linear constraints $\sum_{1\leq i\leq m} y_i = 1$ and $y \geq 0$
  - It makes no difference whether we constrain β to be nonnegative, since the nonnegativity of the a_{ij}'s implies that β is nonnegative in any optimal solution

- Like the minimax LP, the maximin LP is feasible and has a finite optimal value for the objective function

The Dual of the Minimax LP

- Recall that an LP of the form "maximize $c^T x$ subject to $Ax \leq b$ and $x \geq 0$" has as its dual the LP "minimize $y^T b$ subject to $A^T y \geq c$ and $y \geq 0$"

- By putting the column player LP into this standard form, we can mechanically write out the dual of the column player LP

The Dual of the Minimax LP

- We obtain the following dual LP with nonnegative variables $y_1, \ldots, y_m, \beta', \beta''$: Minimize $\beta' - \beta''$ subject to

$$\left(\sum_{1 \le i \le m} y_i \cdot a_{ij} \right) + \beta' - \beta'' \ge 0$$

  for each column j and

$$\sum_{1 \le i \le m} y_i \le 1$$

- Note that this LP is extremely similar to the row player's maximin LP
- We can make it more similar by eliminating the nonnegative variables β' and β'' in favor of a single unrestricted variable β
  - Replace each occurrence of $\beta'' - \beta'$ with β

The Dual of the Minimax LP

- The objective of the dual of the minimax LP is "minimize $-\beta$"
  - Note that this is equivalent to "maximize β", the objective of the row player LP
- The only remaining difference between the dual of the column player LP and the row player LP is that the former includes the constraint $\sum_{1 \leq i \leq m} y_i \leq 1$, but not the stronger constraint $\sum_{1 \leq i \leq m} y_i = 1$
- But since the a_{ij}'s are all nonnegative, it is clear that there is an optimal solution to the dual of the column player LP for which $\sum_{1 \leq i \leq m} y_i = 1$
- In other words, we can add the constraint $\sum_{1 \leq i \leq m} y_i \geq 1$ to the dual of the column player LP without changing the value of an optimal solution

Von Neumann's Minimax Theorem

- Let I, I', and I'' denote the minimax LP (i.e., the column player LP), the maximin LP (i.e., the row player LP), and the dual of the minimax LP, respectively
- Let v, v', and v'' denote the optimal value of the objective function of I, I', and I'', respectively
- From the foregoing discussion, $v' = v''$
- By the strong duality theorem, $v = v''$
- Thus $v = v'$

Discussion of the Minimax Theorem

- In other words, if the colum and row players employ optimal mixed strategies, the payoff to the row player is equal to both
  - The minimax payoff α, as determined by solving the column player's LP to determine an optimal mixed strategy x^*
  - The maximin payoff β, as determined in the row player's LP to determine an optimal mixed strategy y^*
- An interesting consequence is that even if the column player publicly commits to the strategy x^*, the row player is still not incented to deviate from y^*
- Symmetrically, if the row player is known to be using strategy y^*, the column player cannot do better than to play x^*
- In this sense the optimal row and column player solutions together form a stable optimal solution to the given zero-sum game

Remarks on General Bimatrix Games

- Nash showed that every bimatrix game admits mixed strategies x and y for the column and row players, respectively, so that neither player is incented to play a different strategy when the other player's strategy is revealed

- Such a pair of strategies (x, y) is referred to as a Nash equilibrium

- In fact Nash, proved the existence of such equilibria even when there are $k > 2$ players
  - Note that there is a natural way to generalize the notion of a bimatrix game to $k > 2$ players

- Even though Nash's result guarantees the existence of such equilibria, no polynomial-time algorithm is known for computing a Nash equilibrium, even for the special case of two players
  - This is a major open problem in complexity theory