Game theory and equilibria

Huy L. Nguyễn

In the course so far, we have typically described problems where the input is given and the algorithm needs to perform some optimization. In many situation, the inputs come from people with economic incentives and they can adapt their behavior depending on the algorithm as well as the behavior of other people. In this note, we make a brief foray into game theory and related algorithmic problems.

1 Zero-sum games

A game for us is a model for the interaction among agents. Each agent has a payoff function and his goal is to maximize the payoff.

A simple class of games is the zero-sum two player games. There are two players: player 1 has m moves and player 2 has n moves. First player one chooses a move i and then player two makes a response j. The outcome is that player 1 pays A_{ij} to player 2. The game can be described completely by an $m \times n$ payoff matrix A.

An example of a zero sum game is rock/paper/scissor. This game is zero sum because the total payoff is zero. This is not true for all games e.g. in war all side might lose.

An equilibrium is a pair of strategy s_1, s_2 such that each is the best response to the other.

Suppose player 1 plays first with move i_0. To maximize payoff, player 2 will choose $j_0 = \mathrm{argmax}_j A_{i_0,j}$. However, if player 2 play first with move j_0 then player 1 will response with move $i_1 = \mathrm{argmax}_i A_{i,j_0}$. In general, $i_0 \neq i_1$ and there is no equilibrium. One can observe this fact from the example of rock/paper/scissor.

von Neumann realized that if instead of restricting the strategy to play only one move and allowing strategy that play a random move according to some probability distribution, a so-called *mixed* strategy, there always exists an equilibrium.

Suppose player 1 plays move i with probability x_i. We have $x_i \geq 0\ \forall i$ and $\sum_i x_i = 1$. Player 2 plays move j with probability y_j. We have $y_j \geq 0\ \forall j$ and $\sum_j y_j = 1$. The expected payoff for player 2 is $x^T A y$.

If player 1 plays first then the expected payoff for player 2 is $\min_x \max_y x^T A y$. On the other hand, if player 2 plays first then the expected payoff is $\max_y \min_x x^T A y$. von Neumann's theorem states that these two amounts are the same and thus, neither player has an incentive to change their strategy after seeing the other's strategy.

Theorem 1.1. $\min_x \max_y x^T A y = \max_y \min_x x^T A y$

It turns out that this theorem is a consequence of linear programming duality.

2 Non-zero sum games

Now we consider the case where the total payoff is not zero. Now we describe the outcome using two matrices A and B. If payer 1 plays move i and player 2 plays move j then player 1 gets A_{ij}

and player 2 gets B_{ij}. A Nash equilibrium is a pair of strategies, one for each player, such that each one is the optimal response to the other i.e. neither has an incentive to change their strategy after seeing the other's strategy.

One common example is prisoner's dilemma. There are two prisoners who are put in separate rooms and questioned separately. They have two choices: to betray the other or to stay silent.

Prisoner One \ Prisoner Two	silent	betray
silent	Each gets 3 years	One gets 5 years, Two is free
betray	One is free, Two gets 5 years	Each gets 4 years

Regardless of what the other prisoner is doing, each of them has an incentive to betray because they always get a better outcome. Thus, the equilibrium is that both of them betray the other. This is much worse than the socially optimal strategy of both of them staying silent.

Another common example is the game of chicken. Two drivers driving their cars at high speed toward crashing into each other. Each of them has two choices: swerve to avoid the collision (chicken) or go straight on (dare). If they both play dare then they risk injury. If both play chicken then they both live. If one play dare and one play chicken then the one who plays chicken looks bad and the one who play dare gains reputation.

Driver 1 \ Driver 2	chicken	dare
chicken	4/4	1/5
dare	5/1	0/0

There are two equilibria $(chicken, dare)$ and $(dare, chicken)$. Note that the socially optimal $(chicken, chicken)$ is not an equilibrium.

Nash shows that for any payoff matrices A, B, there exists a mixed equilibrium.

3 Approximate Nash equilibrium

Consider a game with all payoffs in the range $[-1, 1]$.

Definition 3.1. *A pair of strategies x, y is an ε-approximate Nash equilibrium if*

- *x is an ε-approximate best response to y i.e. $x^T Ay \geq \max_i \langle A_i, y\rangle - \varepsilon$*
- *y is an ε-approximate best response to x i.e. $y^T B^T x \geq \max_j \langle B_j^T, x\rangle - \varepsilon$*

Theorem 3.2 (Lipton-Markakis-Mehta). *Consider a two player game with $n \times n$ payoff matrices A, B. There is an ε approximate Nash equilibrium where each strategy is an uniform sample from a multi-set of size $O(\log n/\varepsilon^2)$.*

Proof. Consider a Nash equilibrium with strategies x and y. Suppose X, Y are multisets of k independent random samples from the strategies x and y, respectively. Let $\hat{x}_i$ be the fraction of times strategy i is sampled for the row player and similarly for $\hat{y}_j$. We will show that for sufficiently large k, the following conditions hold simultaneously with positive probability:

- For all i, the payoff for row player with strategy i is almost the same whether the column player plays y or $\hat{y}$ i.e. $|(Ay)_i - (A\hat{y})_i| \leq \varepsilon$
- For all j, the payoff for column player with strategy j is almost the same whether the row player plays x or $\hat{x}$ i.e. $|(x^T B)_j - (\hat{x}^T B)_j| \leq \varepsilon$
- $|xAy - \hat{x}Ay| \leq \varepsilon$

- $|xBy - xB\hat{y}| \leq \varepsilon$

When the first condition holds, the payoff for the row player with any strategy is almost the same against y and $\hat{y}$. Because x is the best response for y i.e. $x^T Ay = \max_i \langle A_i, y\rangle$, we conclude that

$$\begin{aligned}
\hat{x}^T A\hat{y} &\geq \hat{x}^T Ay - \varepsilon \quad (\text{because } |(Ay)_i - (A\hat{y})_i| \leq \varepsilon) \\
&\geq xAy - 2\varepsilon \\
&\geq \max_i \langle A_i, y\rangle - 2\varepsilon \\
&\geq \max_i \langle A_i, \hat{y}\rangle - 3\varepsilon \quad (\text{because } |(Ay)_i - (A\hat{y})_i| \leq \varepsilon)
\end{aligned}$$

Similar reasoning works for the column player. To prove the first condition, notice that $\mathbb{E}(A\hat{y})_i = (Ay)_i$. By the Chernoff bound, with $k = O(\log n/\varepsilon^2)$, we have $|(Ay)_i - (A\hat{y})_i| \leq \varepsilon$ with probability at least $1 - 1/n^2$. Thus, by the union bound over all choices of i, the first condition holds. A similar reasoning works for the other conditions. $\square$

This theorem implies an algorithm for finding an approximate Nash equilibrium: try all possible strategies of the above form. This algorithm takes time $n^{O(\log n/\varepsilon^2)}$.