COMPSCI 240: Reasoning Under Uncertainty

Nic Herndon and Andrew Lan

University of Massachusetts at Amherst

Spring 2019

Lecture 21: Game Theory I

Game Theory

- Our textbook does not include the materials related to Game Theory (this & next two lectures).
- We will use lecture slides for the next three lectures to provide a brief introduction to Game Theory.

Game Theory

- Game theory studies what happens when self-interested agents interact.
- Each player (or agent) has his/her own description of which states of the world he/she likes (or the states that benefits the agent the most).
- A player's benefits can be represented using a **payoff matrix** that maps the *states* to *real numbers*.
- For now, let us assume that all players select a single action and play – all players have pure strategy.

Prisoner's Dilemma

- Two suspects are being held pending trial for a crime they are alleged to have committed. The prosecutor does not have clear evidence who's the criminal and needs one of suspects to rat out the other.
- Thus, the prosecutor offers each a deal: **"Give evidence against your partner and you'll go free, unless your partner also confesses. If both confess, both get 5 year sentences. If neither confess, both get 1 year sentences. If you don't confess but your partner does, you get 10 years!"**
- Prisoner A and B have two actions: 1) Confess or 2) Stay silent.
- Suppose the prisoners' interests are in minimizing the number of years in jail.

Prisoner's Dilemma

- The payoff matrix looks like the following:

	B Confesses	B Stays Silent
A Confesses	$-5, -5$	$0, -10$
A Stays Silent	$-10, 0$	$-1, -1$

  In each entry, the first number is A's payoff and the second number is B's payoff.
- What would be best action that each prisoner can take without discussing together?
- Given that B confesses, then what action should A take?
- Given that B stays silent, then what action should A take?
- Similarly, given actions of A, what actions shall B take?

Strict Domination

- A strategy s_i of a player P_i strictly dominates another strategy s_i' of the player, if s_i generates a greater payoff than s_i'.
- More formally, let P_i represent a player i and P_{-i} represent all other players but i. Furthermore, Let s_i and s_i' be two strategies of P_i, and S_{-i} be the set of all strategies of the remaining players.
- Then, s_i strictly dominates s_i' if for all $s_{-i} \in S_{-i}$, it is the case that $u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$, where u_i is the payoff function.

Prisoner's Dilemma

- Confessing strictly dominates all other strategies (stay silent).

	B Confesses	B Stays silent
A Confesses	$-5, -5$	$0, -10$
A Stays silent	$-10, 0$	$-1, -1$

- We have four possible states in this game but there's only one sensible outcome: i.e., both suspects confess.
- It is interesting because staying quiet produces an outcome that is mutually (and individually) better for both suspects.
- This is because each player (suspects) plays the game for their own interests.

Dominated Strategy

- A strategy s_i of a player P_i is **strictly dominated** if some (not all) other strategies s'_i strictly dominates s_i.

Iterated Elimination of Strictly Dominated Strategies (IESDS)

- Intuitively, all strategies that are strictly dominated by other strategies can be ignored, since they can never be best responses to any moves by the other players.
- Consider the following payoff matrix:

	L	C	R
U	$13, 3$	$1, 4$	$7, 3$
M	$4, 1,$	$3, 3$	$6, 2$
D	$-1, 9$	$2, 8$	$8, -1$

- Strategy R of Player 2 is strictly dominated by C, so remove it.

Iterated Elimination of Strictly Dominated Strategies

- The payoff matrix reduced after removing the dominated strategy R.

	L	C
U	$13, 3$	$1, 4$
M	$4, 1,$	$3, 3$
D	$-1, 9$	$2, 8$

- Note that D is strictly dominated by M, so remove it.

Iterated Elimination of Strictly Dominated Strategies

- The pay-off matrix reduced to

	L	C
U	$13, 3$	$1, 4$
M	$4, 1,$	$3, 3$

- Now L is dominated by C, so remove it.
- Then, the only reasonable strategy to be played is M for Player 1 and C for Player 2.

Hawks and Doves

- Two birds meet over a piece of food and have to decide whether to act aggressive (**hawkish**) or passive (**dovish**)
  - If a hawk meets a dove, the hawk gets the food worth 50 points
  - If two hawks meet, they both loose -25 points
  - If two doves meet, they both get 15 points
- Can represent this as:

	B is a Hawk	B is a Dove
A is a Hawk	$-25, -25$	$50, 0$
A is a Dove	$0, 50$	$15, 15$

- We cannot solve this game by looking for strictly dominating strategies.
- We are going to introduce a new way to solve this game.

Nash Equilibrium

Definition

A Nash Equilibrium is a set of strategies for *each player* where no change by one player alone can improve his/her outcome; each player has no incentive to change his/her rather stable strategy.

- We are considering Nash Equilibrium where players do not randomize between two or more strategies (called Pure Strategy Nash Equilibrium).
- We only care about alternating strategies in an individual level, not in a group level where everyone collectively change strategies toward a single strategy.
- A Nash Equilibrium is a law that no one would want to break, even in the absence of an effective police force.
- There may exist multiple Nash Equilibria.

Theorem (Nash)

Every game where each player has a finite number of options, has at least one Nash equilibrium.

Iterated Elimination of Strictly Dominated Strategies vs. Nash Equilibrium

- Strictly dominated strategies cannot be a part of a Nash equilibrium.
- After completing the IESDS, if there exists only one strategy for each player remaining, that strategy set is the unique Nash equilibrium.
- Even if there exists no solutions from the IESDS, there may exist Nash Equilibria.

Prisoner's Dilemma

	B Confesses	B Stays silent
A Confesses	$-5, -5$	$0, -10$
A Stays silent	$-10, 0$	$-1, -1$

- The solution of the IESDS yield that A Confesses and B Confesses.
- Note that this is also a Nash Equilibrium

Stop Light Example

	Go	Stop
Go	$-5, -5$	$1, 0$
Stop	$0, 1$	$-1, -1$

- This has two Nash Equilibria
  - A = Go and B = Stop (e.g., when A has a green light)
  - A = Stop and B = Go (e.g., when B has a green light)
- No Nash Equilibrium where both players play the same pure strategy.