Lecture 15: Game Theory
CS486/686 Intro to Artificial Intelligence

Pascal Poupart
David R. Cheriton School of Computer Science
CIFAR AI Chair at Vector Institute

UNIVERSITY OF WATERLOO

Outline

- Game Theory
- Normal form games
  - Strictly dominated strategies
  - Pure strategy Nash equilibria
  - Mixed Nash equilibria

Multi-agent Decision Making

- Sequential Decision Making
  - Markov Decision Processes
  - Reinforcement Learning
  - Multi-Armed Bandits

- All in single agent environments

- Real world environments: usually more than one agent?
  - Each agent needs to account for other agents' actions/behaviours

Reinforcement Learning

Single agent

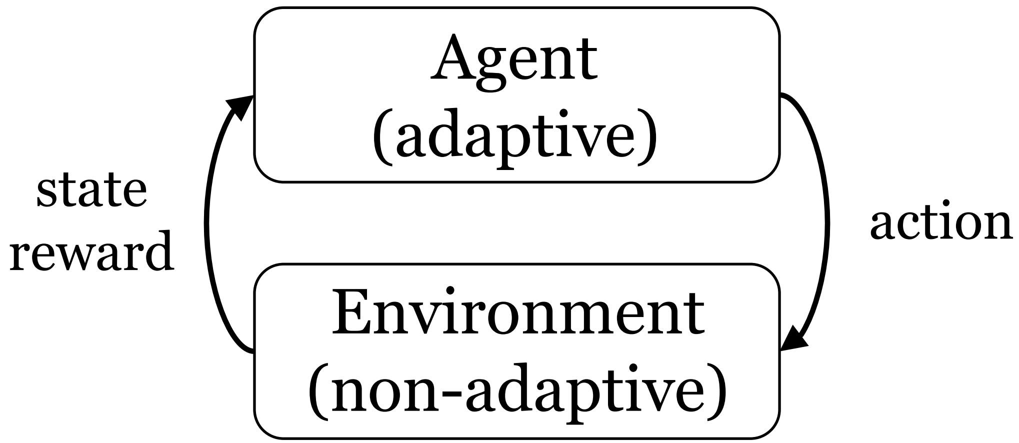

Assumption: stationary transition function

$P_t(s'|s,a) = P_{t'}(s'|s,a)\ \forall t, t'$

Multiple agents

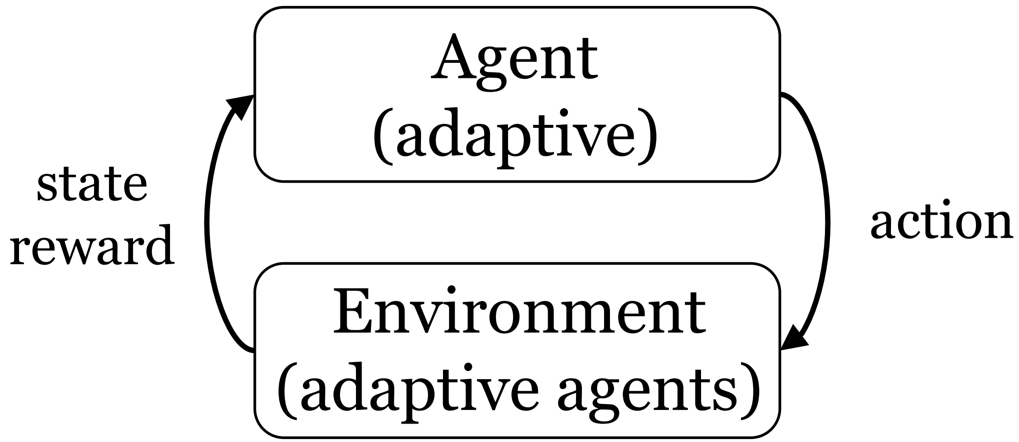

Non-stationary transition function

$P_t(s'|s,a) \neq P_{t'}(s'|s,a)$

Game

- **Game**: Any scenario where outcomes depend on actions of two or more rational and self-interested players
  - **Players** (Decision Makers)
    - Agents within the game (observe states and take actions)
  - **Rational**
    - Agents choose their best actions (unless exploring)
  - **Self-interested**
    - Only care about their own benefits
    - May/May not harm others

Which of these are games?

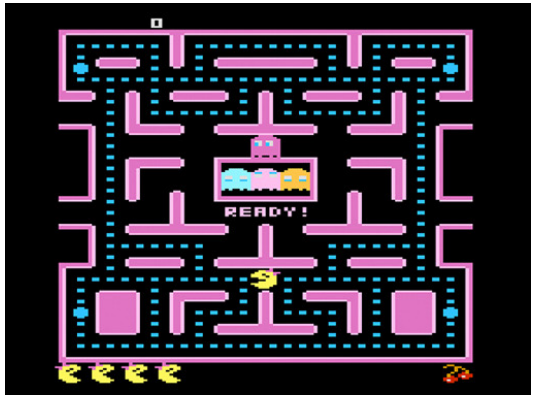

Atari

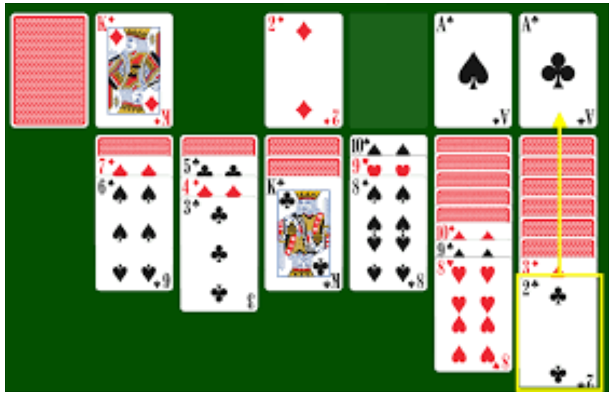

Solitaire

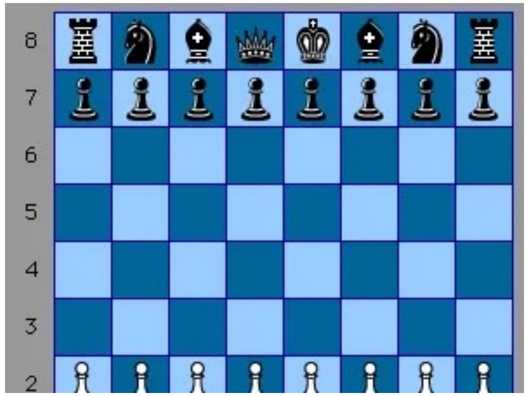

Chess

Go

Game Theory

- **Game Theory**: Mathematical model of strategic interactions among multiple rational agents in a game
  - **Interaction:**
    - One agent directly affects other agent(s)
    - Reward for one agent depends on other agent(s)
  - **Strategic:**
    - Agents maximize their reward by taking into account their influence (through actions) on the game
  - **Multiple:**
    - At-least two agents

Game Theory Applications

- Auctions
- Diplomacy
- Negotiations
- Sports analytics
- Autonomous Driving
- Conversational agents

Categorization of Games

- Games can be
  - **Cooperative:** agents have a common goal
  - **Competitive:** agents have conflicting goals
  - **Mixed:** in between cooperative and competitive
    (agents have different goals, but they are not conflicting)

Cooperative

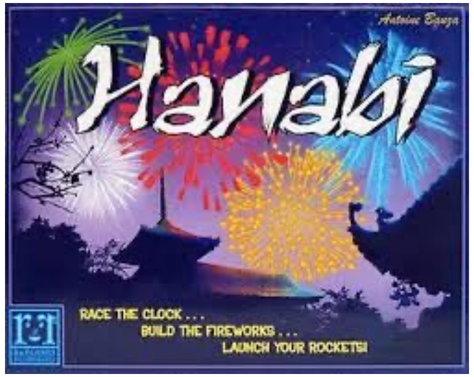

Competitive

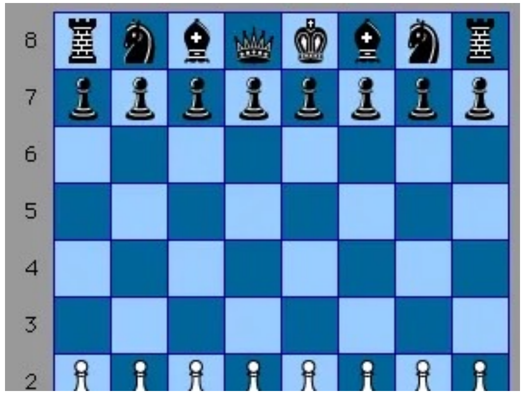

Mixed

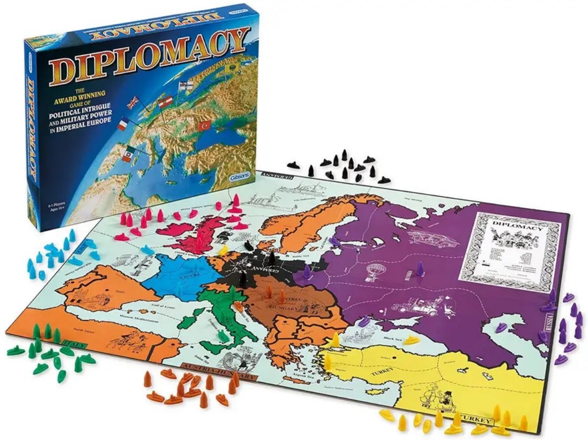

Normal Form Games

- Set of **agents**: $I = 1,2, \dots, N$, where $N \geq 2$
- Set of **actions** for each agent: $A_i = \{a_i^1, \dots, a_i^m\}$
  - Game outcome is a **strategy profile (joint action)**: $\boldsymbol{a} = (a_1, \dots, a_n)$
  - Total space of joint actions: $\boldsymbol{a} \in \{A_1 \times A_2 \times \cdots \times A_N\}$
- **Reward function** for each agent: $R_i: \boldsymbol{A} \to \Re$, where $\boldsymbol{A} = \{A_1 \times A_2 \times \cdots \times A_N\}$

- No state
- Horizon: $h = 1$

Example: Even or Odd

		Agent 2	
		One	Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

R_i

$R_1(a_1, a_2)$

$R_2(a_1, a_2)$

Zero-sum game:

$$\sum_{i=1}^{n} R_i(a_1, \ldots, a_n) = 0$$

$I = \{1,2\}$

$A_i = \{One, Two\}$

An outcome is (One, Two)

$R_1(One, Two) = -3$ and $R_2(One, Two) = 3$

Examples of strategic games

Baseball or Soccer

	B	S
B	2,1	0,0
S	0,0	1,2


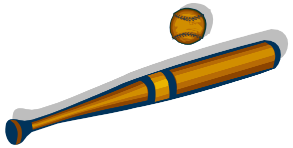

Coordination Game

Chicken

	Cross	Turn
Cross	-1,-1	10,0
Turn	0,10	5,5

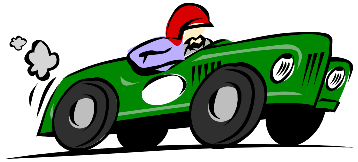


Anti-Coordination Game

Example: Prisoner's Dilemma



	Confess	Don't Confess
Confess	-5,-5	0,-10
Don't Confess	-10,0	-1,-1

Playing a game

- We now know how to describe a game
- Next step – Playing the game!

- Recall, agents are **rational**
  - Let p_i be agent i's beliefs about what its opponents will do
  - Agent i is rational if it chooses to play a_i^* where

$$a_i^* = argmax_{a_i} \sum_{a_{-i}} R(a_i, a_{-i}) p_i(a_{-i})$$

Notation: $a_{-i} = (a_1, \ldots, a_{i-1,} a_{i+1}, \ldots, a_n)$

Dominated Strategies

- **Definition**: A strategy a_i is *strictly dominated* if

$$\exists a_i', \forall a_{-i}, R_i(a_i, a_{-i}) < R(a_i', a_{-i})$$

- A rational agent will never play a strictly dominated strategy!
  - This allows us to solve some games!

Example: Prisoner's Dilemma

	Confess	Don't Confess
Confess	-5,-5	0,-10
~~Don't Confess~~	~~-10,0~~	~~-1,-1~~

	Confess	Don't Confess
Confess	-5,-5	~~0,-10~~

	Confess
Confess	-5,-5

Equilibrium Outcome

Strict Dominance does not capture the whole picture

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	~~6,6~~ 2,2

What strict dominance eliminations can we do?

Nash Equilibrium

- Sometimes an agent's best-response depends on the strategies other agents are playing

- A strategy profile, $\boldsymbol{a}^*$, is a **Nash equilibrium** if no agent has incentive to deviate from its strategy *given that others do not deviate*:

$$\forall i, a_i, \qquad R_i(a_i^*, a_{-i}^*) \geq R_i(a_i, a_{-i}^*)$$

Nash Equilibrium

- Equivalently, $\boldsymbol{a}^*$ is a Nash equilibrium iff $\forall i\ a_i^* = argmax_{a_i} R_i(a_i, a_{-i}^*)$

	A	B	C
A	0,4	4,0	5,3
B	4,0	0,4	5,3
C	3,5	3,5	6,6

(C,C) is a Nash equilibrium because:

$$R_1(C,C) = \max\{R_1(A,C), R_1(B,C), R_1(C,C)\}$$

AND

$$R_2(C,C) = \max\{R_2(C,A), R_2(C,B), R_2(C,C)\}$$

Exercise 1

What are the Nash Equilibria?

	B	S
B	2,1	0,0
S	0,0	1,2

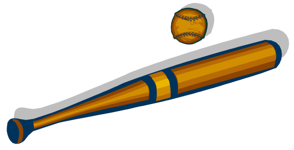

Exercise 2

What are the Nash Equilibria?

		Agent 2: One	Agent 2: Two
Agent 1	One	2,-2	-3,3
	Two	-3,3	4,-4

(Mixed) Nash Equilibria

- Mixed strategy σ_i: σ_i defines a probability distribution over A_i
- Strategy profile: $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_n)$
- Expected utility: $R_i(\boldsymbol{\sigma}) = \sum_a (\prod_j \sigma(a_j)) R_i(a)$
- Nash Equilibrium: σ^* is a (mixed) Nash equilibrium if

$$\forall i\ R_i(\sigma_i^*, \sigma_{-i}^*) \geq R_i(\sigma_i', \sigma_{-i}^*)\ \forall \sigma_i'$$

Finding Mixed Nash Equilibria

- Two players: $\boldsymbol{\sigma} = (\sigma_1, \sigma_2)$
  - Let $p = \sigma_1$ and $q = \sigma_2$
- At the equilibrium:
  - p^* should be best strategy given q^*: $p^* = argmax_p\, R_1(p, q^*)$
  - q^* should be best strategy given p^*: $q^* = argmax_q R_2(p^*, q)$
- Solve system of equations:
  - $\frac{\partial}{\partial p} R_1(p, q) = 0$
  - $\frac{\partial}{\partial q} R_2(p, q) = 0$

Exercise 2 Revisited

		B: One	B: Two
A	One	2,-2	-3,3
	Two	-3,3	4,-4

$$p = \Pr(one)$$
$$q = \Pr(one)$$

How do we determine p and q?

$$R_A(p,q) = 2pq - 3p(1-q) - 3(1-p)q + 4(1-p)(1-q)$$
$$R_B(p,q) = -2pq + 3p(1-q) + 3(1-p)q - 4(1-p)(1-q)$$
$$\frac{\partial}{\partial p} R_A(p,q) = 12q - 7 \rightarrow q = \frac{7}{12}$$
$$\frac{\partial}{\partial q} R_B(p,q) = -12p + 7 \rightarrow p = \frac{7}{12}$$

Exercise 3

$p = Pr(\text{Agent 1 Selects B})$

$q = Pr(\text{Agent 2 Selects B})$

	B	S
B	2,1	0,0
S	0,0	1,2

This game has 3 Nash equilibria (2 pure strategy Nash equilibria and 1 mixed strategy Nash equilibrium). Find them.

$$R_1 = 2.p.q + (1-p)(1-q) \rightarrow \frac{\partial R_1}{\partial p} = 2q + (q-1) = 0$$

$$q = 1/3$$

$$R_2 = p.q + 2(1-p)(1-q)$$

$$\rightarrow \frac{\partial R_2}{\partial q} = p - 2(1-p) = 0 \rightarrow p = 2/3$$

Mixed Nash Equilibrium

- **Theorem** (Nash 1950):

  Every game in which the strategy sets $A_1, \ldots, A_n$ have a finite number of elements has a mixed strategy equilibrium.

John Nash
Nobel Prize in Economics (1994)

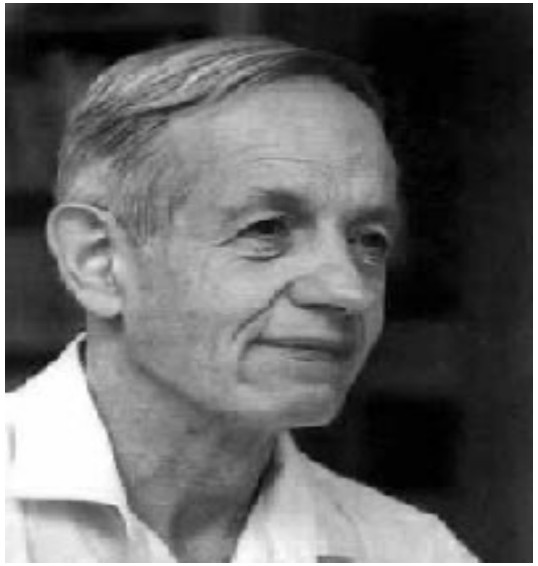

Other Useful Theorems

$a_1^* \leftarrow \sigma_1$ $\quad$ $\sigma_1 \ldots \sigma_n$

- **Theorem:** In an n-player pure strategy game, if iterated elimination of strictly dominated strategies eliminates all but the strategies $(a_1^*,\ldots,a_n^*)$ then these strategies are the unique Nash equilibria of the game

- **Theorem:** Any Nash equilibrium will survive iterated elimination of strictly dominated strategies.

Nash Equilibrium

- Interpretations:
  - Focal points, self-enforcing agreements, stable social convention, consequence of rational inference..

- Criticisms
  - They may not be unique
    - Ways of overcoming this: Refinements of equilibrium concept, Mediation, Learning
  - They may be hard to find
  - People don't always behave based on what equilibria would predict