CS 486/686: Introduction to Artificial Intelligence

Introduction

Introduction

- So far almost everything we have looked at has been in a single-agent setting
  - Today - **Multiagent Decision Making**!
- For participants to act optimally, they must account for how others are going to act
- We want to
  - Understand the ways in which agents interact and behave
  - Design systems so that agents behave the way we would like them to

Hint for the final exam: MAS is my main research area. I like MAS problems. I even enjoy marking MAS questions. There ***will*** be a MAS question on the exam.

Self-Interest

- We will focus on ***self-interested*** MAS
- Self-interested does **not** necessarily mean
  - Agents want to harm others
  - Agents only care about things that benefit themselves
- Self-interested means
  - Agents have their own description of states of the world
  - Agents take actions based on these descriptions

What is Game Theory?

- The study of games!
  - Bluffing in poker
  - What move to make in chess
  - How to play Rock-Paper-Scissors



But also

- auction design
- strategic deterrence
- election laws
- coaching decisions
- routing protocols
- …

What is Game Theory?

- Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**
  - **Group:** Must have more than 1 decision maker
    - Otherwise, you have a decision problem, not a game

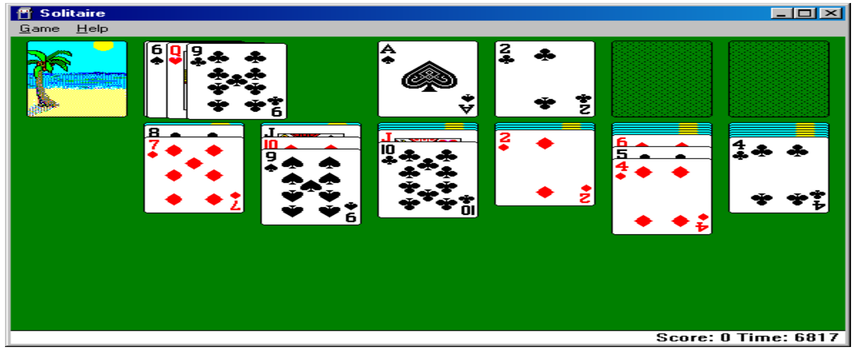

Solitaire is not a game!

What is Game Theory?

- Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents that behave **strategically**

  - **Interaction:** What one agent does directly affects at least one other
  - **Strategic**: Agents take into account that their actions influence the game
  - **Rational**: Agents chose their best actions

Example

- Decision Problem
  - Everyone pays their own bill
- Game
  - Before the meal, everyone decides to split the bill evenly

Strategic Game
(Matrix Game, Normal Form Game)

- Set of agents: I={1,2,.,,,N}
- Set of actions: A_i={a_i^1,…,a_i^m}
- Outcome of a game is defined by a profile a=(a_1,…,a_n)
- Agents have preferences over outcomes
    - Utility functions u_i:A->**R**

Examples

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

I={1,2}
A_i={One,Two}
A_n outcome is (One, Two)
U_1((One,Two))=-3 and U_2((One,Two))=3

Zero-sum game.

$\sum_{i=1}^{n} u_i(o)=0$

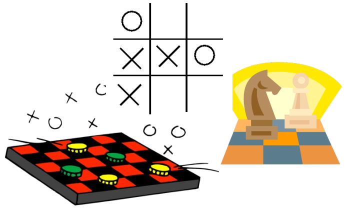

Examples

BoS

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

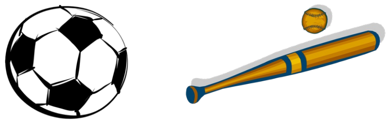

Coordination Game

Chicken

	T	C
T	-1,-1	10,0
C	0,10	5,5

Anti-Coordination Game

Example: Prisoners' Dilemma

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

Playing a Game

- Agents are rational
  - Let p_i be agent *i*'s belief about what its opponents will do
  - **Best response**: $a_i = \text{argmax} \sum_{a_{-i}} u_i(a_i, a_{-i}) p_i(a_{-i})$

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

Dominated Strategies

- a'_i **strictly dominates** strategy a_i if

$$u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i}) \forall a_{-i}$$

- A rational agent will never play a dominated strategy!

Example

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

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

Nash Equilibrium

Key Insight: an agent's best-response depends on the actions of other agents

An action profile a* is a **Nash equilibrium** if no agent has incentive to change given that others do not change

$$\forall i u_i(a_i^*, a_{-i}^*) \geq u_i(a_i', a_{-i}^*) \forall a_i'$$

Nash Equilibrium

Equivalently, a* is a N.E. iff

$$\forall i a_i^* = \arg\max_{a_i} u_i(a_i, a_{-i}^*)$$

	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 N.E. because

$$u_1(C,C) = \max \begin{bmatrix} u_1(A,C) \\ u_1(B,C) \\ u_1(C,C) \end{bmatrix}$$

AND

$$u_2(C,C) = \max \begin{bmatrix} u_2(C,A) \\ u_2(C,B) \\ u_2(C,C) \end{bmatrix}$$

Nash Equilibrium

- If (a_1*,a_2*) is a N.E. then player 1 won't want to change its action given player 2 is playing a_2*

- If (a_1*,a_2*) is a N.E. then player 2 won't want to change its action given player 1 is playing a_1*

-5,-5	0,-10
-10,0	-1,-1

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

Another Example

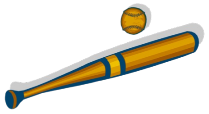

	B	S
B	2,1	0,0
S	0,0	1,2

2 Nash Equilibria

Coordination Game

Yet Another Example

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

(Mixed) Nash Equilibria

- **(Mixed) Strategy**: s_i is a probability distribution over A_i
- **Strategy profile**: $s=(s_1,...,s_n)$
- **Expected utility**: $u_i(s)=\Sigma_a \Pi_j s(a_j) u_i(a)$
- **Nash equilibrium**: s* is a (mixed) Nash equilibrium if

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*) \forall s_i'$$

Yet Another Example

q	One	Two
p One	2,-2	-3,3
p Two	-3,3	4,-4

How do we determine p and q?

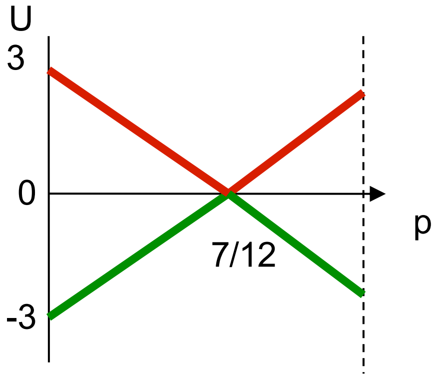

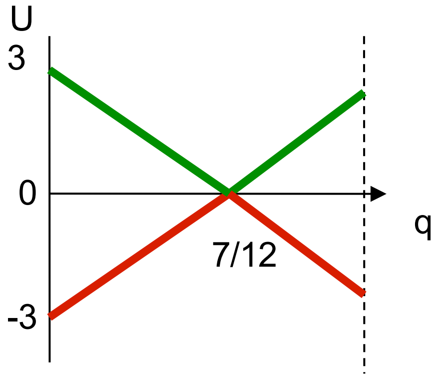

Yet Another Example

		q One	Two
p	One	2,-2	-3,3
	Two	-3,3	4,-4

How do we determine p and q?

Exercise

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

This game has 3 Nash Equilibrium (2 pure strategy NE and 1 mixed strategy NE).

Mixed Nash Equilibrium

Theorem (Nash 1950): Every game in which the action sets are finite, has a mixed strategy equilibrium.

John Nash
Nobel Prize in Economics (1994)

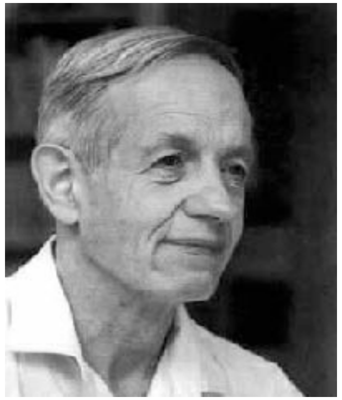

Finding NE

- Existence proof is *non-constructive*
- Finding equilibria?
  - 2 player zero-sum games can be represented as a linear program (polynomial)
  - For arbitrary games, the problem is in PPAD
  - Finding equilibria with certain properties is often NP-hard

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are **habitual** criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

How do we define payoffs?

What is the strategy space?

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are **habitual** criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

…

How do we define payoffs?

Average reward

Discounted Awards

…

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are habitual criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

…

Strategy space becomes significantly larger!

S:H→A where H is the **history** of play so far

Can now reward and punish past behaviour, worry about reputation, establish trust,…

Repeated Games

Recall the Prisonner's Dilemma. What if the prisoners are habitual criminals?

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

-5,-5	0,-10
-10,0	-1,-1

...

Grim Strategy: In first step cooperate. If opponent defects at some point, then defect forever

Tit-for-Tat: In first step cooperate. Copy whatever opponent did in previous stage.

Summary

Definition of a Normal Form Game

Dominant strategies

Nash Equilibria