

Lecture 14

June 16, 2005

CS 486/686

Outline

- Multi-agent systems
- Game theory
- Russell and Norvig: Sect 17.6

Course Topics

- Search
 - Uninformed and heuristic search
 - CSP's and optimization
 - Game playing
- Reasoning under uncertainty
 - Probability theory, utility theory and decision theory
 - Bayesian networks and decision networks
 - Multi-agent systems
- Learning
 - Decision trees, neural networks, ensemble learning, reinforcement learning
- Specialized areas
 - Natural language processing, computational vision and robotics

Multi-agent systems


- So far...
 - Single agent optimizing some objectives in a possibly uncertain environment
 - But, what if there are several agents?
- Multi-agent systems
 - Two (or more) agents can influence the world
 - How should an agent act given that it shares “control” with other agents?

Multi-agent Systems

- Search techniques for deterministic games with alternating play
 - Minimax algorithm
 - Alpha-beta pruning
- Today:
 - Extend decision theory to multi-agent systems
 - View other agents as sources of uncertainty
 - Framework: **Game theory**

What is game theory?

- Game theory is a formal way to analyze **interactions** among a **group** of **rational** agents who 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 who behave **strategically**
 - **Interaction**: What one agent does directly affects at least one other agent in the group
 - **Rational**: An agent chooses its best action
 - **Strategic**: Agents take into account how other agents influence the game

Games

- Examples:
 - Chess, soccer, poker, etc.
 - Elections
 - Auctions, Trades
 - Taxation system
 - Negotiation
 - Packet routing protocols,
 - Driving laws

Two aspects

- **Agent design**
 - Given a game, what is a rational strategy?
 - Ex: playing chess, driving, voting, filling up an income tax report, etc.
- Mechanism design
 - Given that agents behave rationally, what should the rules of the game be?
 - Ex: designing driving laws, an election, a taxation system, an auction, etc.

Strategic Games (aka normal form)

- Formally: $\langle I, \{S_i\}, \{U_i\} \rangle$
- Set of **agents** $I = \{1, 2, \dots, n\}$
- Each agent i can choose a **strategy** $s_i \in S_i$
- Outcome of the game is defined by a **strategy profile** $(s_1, \dots, s_n) \in S$
- Agents have **preferences** over the outcomes
 - utility functions: $U_i(s_1, \dots, s_n) \in \mathfrak{R}$

Example: Election

- **Agents:** electors
- **Strategies:** possible votes for different candidates
- **Outcome:** set of all votes determines a winner (elected candidate)
- **Utility fn:** preferences for each candidate

Simple Games

- Assumptions:
 - Single decision
 - Deterministic game
 - Fully observable game
 - Simultaneous play
- Possible to relax those assumptions... but this is beyond the scope of this course

Example: Even or Odd

Agent 2

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

**Zero-sum
game.**

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

$I = \{1, 2\}$

$S_i = \{\text{One}, \text{Two}\}$

An outcome is (One, Two)


$U_1((\text{One}, \text{Two})) = -3$ and $U_2((\text{One}, \text{Two})) = 3$



Examples of strategic games


Baseball or Soccer

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



Chicken

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



Coordination Game

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 opponent will do
 - Agent i is rational if it chooses to play strategy s_i^* where

$$s_i^* = \operatorname{argmax}_{s_i} \sum_{s_{-i}} u_i(s_i, s_{-i}) p_i(s_{-i})$$

Notation: $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_n)$

Dominated Strategies

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

$$\exists s_i', \forall s_{-i}, u_i(s_i, s_{-i}) < u_i(s_i', s_{-i})$$

- 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

What strict dominance eliminations can we do?

None...

So what should the players of this game do?

Nash Equilibrium

- Sometimes an agent's best-response depends on the strategies other agents are playing
- A strategy profile, s^* , is a **Nash equilibrium** if no agent has incentive to deviate from its strategy *given that others do not deviate*:

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

Nash Equilibrium

- Equivalently, s^* is a N.E. iff
 $\forall i \quad s_i^* = \operatorname{argmax}_{s_i} u_i(s_i, s_{-i}^*)$

	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

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

Another example






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

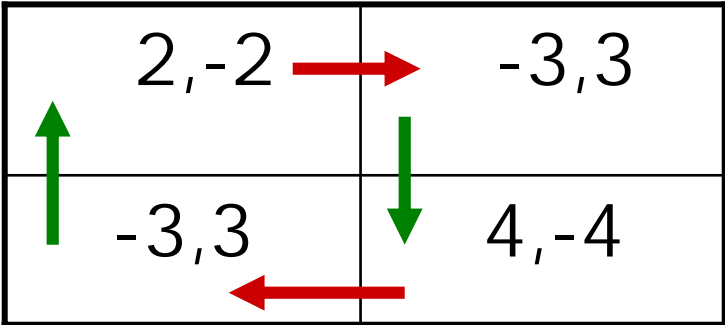
The table shows a coordination game with two players, B and S. The strategies are B and S. The payoffs are (2,1) for (B,B), (0,0) for (B,S), (0,0) for (S,B), and (1,2) for (S,S). Red arrows point from (0,0) to (2,1) and from (0,0) to (1,2). Green circles highlight the cells (2,1) and (1,2), and green arrows point from (0,0) to (2,1) and from (0,0) to (1,2).

2 Nash Equilibria

Coordination Game

Yet another example

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



There is no **PURE** strategy Nash Equilibrium for this game

(Mixed) Nash Equilibria

- **Mixed strategy σ_i :**
 - $\sigma_i \in \Sigma_i$ defines a probability distribution over S_i
- **Strategy profile: $\sigma = (\sigma_1, \dots, \sigma_n)$**
- **Expected utility: $u_i(\sigma) = \sum_{s \in S} (\prod_j \sigma(s_j)) u_i(s)$**
- **Nash Equilibrium: σ^* is a (mixed) Nash equilibrium if**

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

Yet another example

		One	B	Two	
A	One	2, -2		-3, 3	$p = \text{Pr}(\text{one})$ $q = \text{Pr}(\text{two})$
	Two	-3, 3		4, -4	

How do we determine p and q ?

$$U_A(p, q) = 2pq - 3p(1-q) - 3(1-p)q + 4(1-p)(1-q)$$

$$U_B(p, q) = -2pq + 3p(1-q) + 3(1-p)q - 4(1-p)(1-q)$$

$$\frac{\partial}{\partial p} U_A(p, q) = 12q - 7 \Rightarrow q = \frac{7}{12}$$

$$\frac{\partial}{\partial q} U_B(p, q) = -12p + 7 \Rightarrow p = \frac{7}{12}$$

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). Find them.

Mixed Nash Equilibrium

- **Theorem (Nash 50):**

Every game in which the strategy sets, S_1, \dots, S_n have a finite number of elements has a mixed strategy equilibrium.

John Nash
Nobel Prize in Economics (1994)



Other Useful Theorems

- **Thm:** In an n -player pure strategy game $G=(S_1, \dots, S_n; u_1, \dots, u_n)$, if iterated elimination of strictly dominated strategies eliminates all but the strategies (S_1^*, \dots, S_n^*) then these strategies are the unique NE of the game
- **Thm:** Any NE 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 (ultimatum games and notions of fairness,...)

Bayesian Games

- What should player A do?

		Player B	
		L	R
Player A	U	3, ?	-2, ?
	D	0, ?	6, ?

Question: When does such a situation arise?

Bayesian Games

- Hockey lover get 2 units for watching hockey and 1 unit for watching curling
- Curling lover gets 2 units for watching curling and 1 unit for watching hockey
- Pat is a hockey lover
- Pat thinks that Chris is probably a hockey lover, but is not sure

		Chris	
		H	C
Pat	H	2,2	0,0
	C	0,0	1,1

With 2/3 chance

		Chris	
		H	C
Pat	H	2,1	0,0
	C	0,0	1,2

With 1/3 chance

Bayesian Games

- In a Bayesian game each player has a **type**
- All players know their own type, but have only a probability distribution over their opponents' types
- **Game G**
 - Set of action spaces: A_1, \dots, A_n
 - Set of type spaces: T_1, \dots, T_n
 - Set of beliefs: P_1, \dots, P_n
 - Set of payoff functions: u_1, \dots, u_n

 - $P_i(t_{-i} | t_i)$ is the prob distribution of the types for the other players, given player i has type t_i
 - $u_i(a_1, \dots, a_n; t_i)$ is the utility (payoff) to agent i if player j chooses action a_j and agent i has type $t_i \in T_i$

Knowledge Assumptions (Who knows what)

- All players know A_i 's, T_i 's, P_i 's and u_i 's
- The i 'th player knows t_i
but not $t_1, t_2, \dots, t_{i-1}, t_{i+1}, \dots, t_n$
- All players know that all players know the above
- And they know that they know that they know..... (common knowledge)
- **Def:** A **strategy** $s_i(t_i)$ in a Bayesian game is a mapping from T_i to A_i (i.e. it specifies what action should be taken for each type)

Back to our game

- $A_1 = \{H, C\}$ $A_2 = \{H, C\}$
- $T_1 = \{hl, cl\}$ $T_2 = \{hl, cl\}$
- P_1
 - $P_1(t_2=hl | t_1=hl) = 2/3$, $P_1(t_2=cl | t_1=hl) = 1/3$, $P_1(t_2=hl | t_1=cl) = 2/3$,
 $P_1(t_2=cl | t_1=cl) = 1/3$
- P_2
 - $P_2(t_1=hl | t_2=hl) = 1$, $P_2(t_1=cl | t_2=hl) = 0$, $P_2(t_1=hl | t_2=cl) = 1$,
 $P_2(t_1=cl | t_2=cl) = 0$
- U_1
 - $u_1(H, H, hl) = 2$, $u_1(H, H, cl) = 1$, $u_1(H, C, hl) = 0, \dots$
- U_2
 - $u_2(H, H, hl) = 2$, $u_2(H, H, cl) = 1$, $u_2(H, C, cl) = 0, \dots$

Bayesian Nash Equilibrium

- A set of strategies (s_1^*, \dots, s_n^*) are a Pure Bayesian Nash Equilibrium if and only if for each player i , and for all possible types $t_i \in T_i$

$$s_i^*(t_i) = \operatorname{argmax}_{a_i \in A_i} \sum_{t_{-i}} u_i(a_i, s_{-i}^*(t_{-i}))$$

No player, for any of their type, wants to change their strategy

Next Class

- Reinforcement Learning
- Russell and Norvig: Chapter 21