Complex Decision with Multiple Agents

Multiple agents?

- We have concentrated on making decisions in uncertain environments (sthocastic and or partially observed)
- But what if the uncertainty is due to other agents and the decisions they make?
- And what if the decisions of those agents are in turn influenced by our decisions?
  - We are talking about agents that have goals or preferences, not throw of a coin...
- We can model these others agents with utility functions...

Main References for this chapter

- R U S S E L , S . ; N O R V I G , P . Inteligência Artificial: Uma abordagem moderna. 3a. ed. Rio de Janeiro: Elsevier Editora, 2009 (Cap. 17)

- MULTIAGENT SYSTEMS Algorithmic, Game-Theoretic, and Logical Foundations Yoav Shoham and Kevin Leyton-Brown. New york press. 2010.

Utilities and Preferences for Agents

- Assume we have just two agents: $Ag = \{i, j\}$
- Agents are assumed to be *self-interested*: they *have preferences over how the environment is*
- Assume $\Omega = \{\omega_1, \omega_2, \ldots\}$ is the set of "outcomes" that agents have preferences over
- We capture preferences by *utility functions*:

$$u_i = \Omega \to \mathbb{R}$$
$$u_j = \Omega \to \mathbb{R}$$

- Utility functions lead to *preference orderings* over outcomes:

$$\omega' \succ_i \omega \text{ means } u_i(\omega) < u_i(\omega')$$

$$\omega \succ_j \omega' \text{ means } u_j(\omega) > u_j(\omega')$$

Multiagent Encounters

- We need a model of the environment in which these agents will act...
  - agents simultaneously choose an action to perform, and as a result of the actions they select, an outcome in Ω will result
  - the *actual* outcome depends on the *combination* of actions

- Environment behavior may be given by *state transformer function*:

$$\tau : \underbrace{Ac}_{\text{agent } i\text{'s action}} \times \underbrace{Ac}_{\text{agent } j\text{'s action}} \rightarrow \Omega$$

Non-cooperative Game Theory

- What is it?
  - mathematical study of interaction between rational, self-interested agents
- Why is it called non-cooperative?
  - while it's most interested in situations where agents' interests conflict, it's not restricted to these settings
  - the key is that the individual is the basic modeling unit, and that individuals pursue their own interests
    - cooperative/coalitional game theory has teams as the central unit, rather than agents

Game theory has at least two ways

- **Agent design:** We can analyze the agent's decisions and compute the expected utility for each decision (under the assumption that other agents are acting optimally according to utility theory)

- **Mechanism design:** it might be possible to define the rules of the environment so that the **collective good of all agents** is maximized when each agent adopts the game-theoretic solution that maximizes its own utility
  - For example, it can help design the protocols for a collection of Internet traffic routers so that each router has an incentive to act in such a way that global throughput is maximized (Let's see TCP game)

TCP Game

- Internet traffic is governed by the TCP (Transmission Control protocol)
- One feature of TCP is the backoff (*delay*) mechanism; if two computers send information packets into the network and they cause conflict,
  - Both computers back off and reduce the rate for a while until the conflict subsides
- Would it be a good idea to use a defective implementation that do not back off??

TCP Game - 2

- Let's say there are just two computers: Yours and your colleague computer's
- You have two possible strategies: C (for using a correct implementation) and D (for using a defective one)
- If both you and your colleague adopt C then your average packet delay is 1 ms, but if If you both adopt D the delay is 3 ms, because of additional overhead at the network router
- If one of you adopts D and the other adopts C then the D adopter will experience no delay at all, but the C adopter will experience TCP user's a delay of 4 ms
- What should a rational agent do?

TCP Game – 3

- Put it in a table we would have:

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

- If both agents are rational, what they would do?

- That is another version of the problem Prisioner Dilemma. We will talk more about that later...

Defining Games

- Finite, n-person game: $\langle N, A, u \rangle$:
  - N is a finite set of n players, indexed by i
  - $A = A_1 \times \ldots \times A_n$, where A_i is the action set for player i
    - $a \in A$ is an action profile, and so A is the space of action profiles
  - $u = \langle u_1, \ldots, u_n \rangle$, a utility function for each player, where $u_i : A \mapsto \mathbb{R}$

- Writing a 2-player game as a matrix:
  - row player is player 1, column player is player 2
  - rows are actions $a \in A_1$, columns are $a' \in A_2$
  - cells are outcomes, written as a tuple of utility values for each player

Normal (Strategic) Form Games

- Normal Form (Strategic Form): Outcome depends only on agent's actions
- Non-normal form: outcome may depends on environment (randomnly)

Example: Prisioner's dilemma

- The players of the game are two prisoners suspected of a crime
- The prisoners are taken to separate interrogation rooms, and each can either "confess"to the crime or "deny" it (or, alternatively, "cooperate" or "defect")
- Their payoff values can be interpreted as the length of jail term each of prisoner gets in each scenario

Prisioner's dilemma

- Player 2

- Player 1

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

- Shown as (Player 1 utility, Player 2 utility)

Strategies, Dominance and Rationality

- We say strategy s for player p **strongly dominates** strategy s' if the outcome for s is better for p than the outcome for s' , for every choice of strategies by the other player(s)

- **Strategy s weakly** dominates s' if s is better than s on at least one strategy profile and no worse on any other

- A **dominant strategy** is a strategy that dominates all others

- It is **irrational** to play a **dominated** strategy, and **irrational** not to play a **dominant** strategy if one exists

Prisioner's dilemma

- Is there a dominant strategy?

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Prisoner's dilemma

Prisoner's dilemma is any game

	C	D
C	a, a	b, c
D	c, b	d, d

with $c > a > d > b$.

Games of Pure Competition

Players have exactly opposed interests

- There must be precisely two players (otherwise they can't have exactly opposed interests)
- For all action profiles $a \in A$, $u_1(a) + u_2(a) = c$ for some constant c
  - Special case: zero sum

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

- Matchinig pennies: Row player wants equals pennies the other different

Games of Cooperation: Side of Road

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$

	Left	Right
Left	1	0
Right	0	1

Another Example: Rock-paper-scissors

	Rock	Paper	Scissors
Rock	0	−1	1
Paper	1	0	−1
Scissors	−1	1	0

One agents tries to maximize the utility, the other minimize it!

General Games

- The most interesting games combine elements of cooperation and competition

		Husband B	Husband F
Wife	B	2, 1	0, 0
	F	0, 0	1, 2

- Battle of Sex: Husband and wife wish to go to the movies, and they can select among two movies:"Fight" or "Beauty" Movies
- They prefer to go together rather than to separate movies, but while the wife (player 1) prefers B the husband (player 2) prefers F

Analyzing games

- We've defined some canonical games, and thought about how to play them. Now let's examine the games from the outside
- From the point of view of an outside observer, can some outcomes of a game be said to be better than others?
  - we have no way of saying that one agent's interests are more important than another's
  - intuition: imagine trying to find the revenue-maximizing outcome when you don't know what currency has been used to express each agent's payoff
- Are there situations where we can still prefer one outcome to another?

Pareto Optimatility

- Idea: sometimes, one outcome o is at least as good for every agent as another outcome o', and there is some agent who strictly prefers o to o'
  - in this case, it seems reasonable to say that o is better than o'
  - we say that o Pareto-dominates o'.

- An outcome o^* is Pareto-optimal if there is no other outcome that Pareto-dominates it.
  - can a game have more than one Pareto-optimal outcome?
  - does every game have at least one Pareto-optimal outcome?

Pareto Optimatillity in Side of the road?

	Left	Right
Left	1	0
Right	0	1

	Left	Right
Left	1	0
Right	0	1

- An outcome is Pareto Optimal if There is NO other outcome that Pareto dominates it
- Outocome o Pareto-dominates o' if it is at least as good for every agent and there is some agent who strictly prefers o 'to o
- **So, We can also say that o is Pareto Optimal if there is no other outcome where no agents is worse off and at least one strictly prefers it**

Pareto Optimatillity in Battle of Sex

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- An outcome is Pareto Optimal if There is NO other outcome that Pareto dominates it

- Outocome o Pareto-dominates o' if it is at least as good for every agent and there is some agent who strictly prefers o 'to o

Pareto Optimatillity in Prisioner Dillema

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

- An outcome is Pareto Optimal if There is NO other outcome that Pareto dominates it

- Outocome o Pareto-dominates o' if it is at least as good for every agent and there is some agent who strictly prefers o 'to o

Pareto Optimatility and Prisioner's Dilemma

The Prisoner's Dilemma

- (*C*,*C*) is Pareto optimal
  - No profile gives both players a higher payoff
- (*D*,*C*) is Pareto optimal
  - No profile gives player 1 a higher payoff
- (*D*,*C*) is Pareto optimal - same argument
- (*D*,*D*) is Pareto dominated by (*C*,*C*)
  - But ironically, (*D*,*D*) is the dominant strategy equilibrium

- We will see more dominant strategies soon...

Pareto Optimatillity in Matching Pennies

	Heads	Tails
Heads	1,−1	−1, 1
Tails	−1, 1	1,−1

	Heads	Tails
Heads	1,−1	−1, 1
Tails	−1, 1	1,−1

?

- An outcome is Pareto Optimal if There is NO other outcome that Pareto dominates it
- Outocome o Pareto-dominates o' if it is at least as good for every agent and there is some agent who strictly prefers o 'to o

Pareto Optimatility in Examples

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

Pareto Optimatility in Examples

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

Best Response and Nash Equilibrium

- If you knew what everyone else was going to do, it would be easy to pick your own action
- Let $a_{-i} = \langle a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n \rangle$.
  - now $a = (a_{-i}, a_i)$

- Best response: $a_i^* \in BR(a_{-i})$ iff $\forall a_i \in A_i,\ u_i(a_i^*, a_{-i}) \geq u_i(a_i, a_{-i})$

Nash Equilibrium

- Now let's return to the setting where no agent knows anything about what the others will do
- What can we say about which actions will occur?

- Idea: look for stable action profiles.
- $a = \langle a_1, \ldots, a_n \rangle$ is a ("pure strategy") Nash equilibrium iff $\forall i,\ a_i \in BR(a_{-i})$.

Nash Equilibrium in Examples

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

	Left	Right
Left	1	0
Right	0	1

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

Nash Equilibria in Examples

	C	*D*
C	−1, −1	−4, 0
D	0, −4	−3, −3

	Left	Right
Left	1	0
Right	0	1

?

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

?

	Heads	Tails
Heads	1, −1	−1, 1
Tails	−1, 1	1, −1

?

- Prisoner Dilemma: the Nash equilibrium is the only non-Pareto-optimal outcome!!
- The others have no [pure] Nash Equlibria

Mixed Strategies: Prob. Distribution over actions

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing randomly
- Define a strategy s_i for agent i as any probability distribution over the actions A_i.
  - pure strategy: only one action is played with positive probability
  - mixed strategy: more than one action is played with positive probability
    - these actions are called the support of the mixed strategy
- Let the set of all strategies for i be S_i
- Let the set of all strategy profiles be $S = S_1 \times \ldots \times S_n$.

Utility under Mixed Strategies

- What is your payoff if all the players follow mixed strategy profile $s \in S$?
  - We can't just read this number from the game matrix anymore: we won't always end up in the same cell
- Instead, use the idea of expected utility from decision theory:

$$u_i(s) = \sum_{a \in A} u_i(a) Pr(a|s)$$

$$Pr(a|s) = \prod_{j \in N} s_j(a_j)$$

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Matching Pennies's Nash Equilibrium

Agent 2 / Agent 1	Heads	Tails
Heads	1, –1	–1, 1
Tails	–1, 1	1, –1

- Each agent has a penny
- Each agent independently chooses to display his/her penny heads up or tails up
- Easy to see that in this game, no pure strategy could be part of a Nash equilibrium
  - For each combination of pure strategies, one of the agents can do better by changing his/her strategy
    - for (Heads,Heads), agent 2 can do better by switching to Tails
    - for (Heads,Tails), agent 1 can do better by switching to Tails
    - for (Tails,Tails), agent 2 can do better by switching to Heads
    - for (Tails,Heads), agent 1 can do better by switching to Heads
- But there's a mixed-strategy equilibrium:
  - (s,s), where $s(\text{Heads}) = s(\text{Tails}) = ½$

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- It's hard in general to compute Nash equilibria, but it's easy when you can guess the support
- For BoS, let's look for an equilibrium where all actions are part of the support

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Let's player 2 play B with probability p, F with 1-p
- If player 1 best-responses with a mixed strategy, player 2 must make him indefferent B and F (why?)
  - Otherwise, player 1 would be better off choosing a pure strategy, according to which she only played the better of her actions

$$u_1(B) = u_1(F)$$

$$2p + 0(1-p) = 0p + 1(1-p)$$

$$p = \frac{1}{3}$$

Computing Mixed Strategy: Battle of Sexes

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

- Likewise, player 1 must randomize to make player 2 indifferent.
  - Why is player 1 willing to randomize?
- Let player 1 play B with q, F with $1-q$.

$$u_2(B) = u_2(F)$$
$$q + 0(1-q) = 0q + 2(1-q)$$
$$q = \frac{2}{3}$$

- Thus the mixed strategies $(\frac{2}{3}, \frac{1}{3})$, $(\frac{1}{3}, \frac{2}{3})$ are a Nash equilibrium.

Interpreting Mixed Strategies

What does it mean to play a mixed strategy? Different interpretations:

- Randomize to confuse your opponent
  - consider the matching pennies example
- Players randomize when they are uncertain about the other's action
  - consider battle of the sexes
- Mixed strategies are a concise description of what might happen in repeated play: count of pure strategies in the limit
- Mixed strategies describe population dynamics: 2 agents chosen from a population, all having deterministic strategies. MS is the probability of getting an agent who will play one PS or another.

Summary

- Computing Nash equilibira, Pareto Optimals, dominant strategies are relevant to game theory and also to create agents that have to act in strategic environments

- Game theory is very relevant to economics, but also for builing **mulitagent sytems** or **automated mechanism design:** two active topics of research in AI

- **Multiagent systems** are those systems that include multiple autonomous entities with either diverging information or diverging interests, or both

- **Automated mechanism design**: where the rules of the game, i.e. the mecanism is automatically created for the setting and objective at hand (for instance, some function that describes collective good of all agents is maximized or some goals is searched for