Game Theory Intro

Lecture 3

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. Example Matrix Games

Self-interested agents

- What does it mean to say that an agent is self-interested?
  - not that they want to harm other agents
  - not that they only care about things that benefit them
  - that the agent has its own description of states of the world that it likes, and that its actions are motivated by this description

Self-interested agents

- What does it mean to say that an agent is self-interested?
  - not that they want to harm other agents
  - not that they only care about things that benefit them
  - that the agent has its own description of states of the world that it likes, and that its actions are motivated by this description
- We capture this by saying that each agent has a utility function: a mapping from states of the world to real numbers, indicating level of happiness with that state of the world
  - quantifies degree of preference across alternatives
  - allows us to understand the impact of uncertainty on these preferences
  - Decision-theoretic rationality: take actions to maximize expected utility.

Why Utility?

- Why would anyone argue with the idea that an agent's preferences could be described using a utility function?

Why Utility?

- Why would anyone argue with the idea that an agent's preferences could be described using a utility function?
  - why should a single-dimensional function be enough to explain preferences over an arbitrarily complicated set of alternatives?
  - Why should an agent's response to uncertainty be captured purely by the *expected value* of his utility function?
- It turns out that the claim that an agent has a utility function is substantive.
- There's a famous theorem (von Neumann & Morgenstern, 1944) that derives the existence of a utility function from a more basic preference ordering and axioms on such orderings.
  - see Theorem 3.1.18 in the book, which includes a proof.

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. Example Matrix Games

Non-Cooperative Game Theory

- What is it?

Non-Cooperative Game Theory

- What is it?
  - mathematical study of interaction between rational, self-interested agents

Non-Cooperative Game Theory

- What is it?
  - mathematical study of interaction between rational, self-interested agents
- Why is it called non-cooperative?

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

TCP Backoff Game

TCP Backoff Game

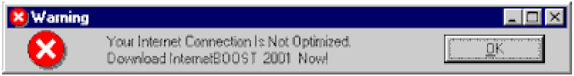

Should you send your packets using correctly-implemented TCP (which has a "backoff" mechanism) or using a defective implementation (which doesn't)?

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.

TCP Backoff Game

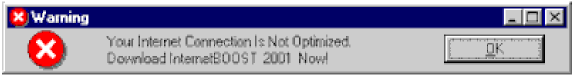

Should you send your packets using correctly-implemented TCP (which has a "backoff" mechanism) or using a defective implementation (which doesn't)?

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.
- Play this game with someone near you. Then find a new partner and play again. Play five times in total.

TCP Backoff Game

- Consider this situation as a two-player game:
  - both use a correct implementation: both get 1 ms delay
  - one correct, one defective: 4 ms delay for correct, 0 ms for defective
  - both defective: both get a 3 ms delay.
- Questions:
  - What action should a player of the game take?
  - Would all users behave the same in this scenario?
  - What global patterns of behaviour should the system designer expect?
  - Under what changes to the delay numbers would behavior be the same?
  - What effect would communication have?
  - Repetitions? (finite? infinite?)
  - Does it matter if I believe that my opponent is rational?

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

Games in Matrix Form

Here's the TCP Backoff Game written as a matrix ("normal form").

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

Lecture Overview

1. Self-interested agents
2. What is Game Theory?
3. Example Matrix Games

More General Form

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
- Thus, we only need to store a utility function for one player
  - in a sense, it's a one-player game

Matching Pennies

One player wants to match; the other wants to mismatch.

	Heads	Tails
Heads	1	-1
Tails	-1	1

Matching Pennies

One player wants to match; the other wants to mismatch.

	Heads	Tails
Heads	1	-1
Tails	-1	1

Play this game with someone near you, repeating five times.

Rock-Paper-Scissors

Generalized matching pennies.

	Rock	Paper	Scissors
Rock	0	-1	1
Paper	1	0	-1
Scissors	-1	1	0

...Believe it or not, there's an annual international competition for this game!

Games of Cooperation

Players have exactly the same interests.

- no conflict: all players want the same things
- $\forall a \in A, \forall i, j,\ u_i(a) = u_j(a)$
- we often write such games with a single payoff per cell
- why are such games "noncooperative"?

Coordination Game

Which side of the road should you drive on?

	Left	Right
Left	1	0
Right	0	1

Coordination Game

Which side of the road should you drive on?

	Left	Right
Left	1	0
Right	0	1

Play this game with someone near you. Then find a new partner and play again. Play five times in total.

General Games: Battle of the Sexes

The most interesting games combine elements of cooperation *and* competition.

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

General Games: Battle of the Sexes

The most interesting games combine elements of cooperation *and* competition.

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

Play this game with someone near you. Then find a new partner and play again. Play five times in total.