CS-E4800 Artificial Intelligence

Jussi Rintanen

Department of Computer Science
Aalto University

March 9, 2017

Centralized versus Distributed Systems

Centralized AI with Distributed Components

Multi-actor systems (fleet of robots, vehicles, machines) under central control and full communication work like single-agent systems (lecture until now!)

Distributed AI and Multi-Agent Systems

- Collaboration of independent agents
- Competition of independent agents

(Independence = Agents have different "owners")

Autonomous Uber Taxi Cab Service

Exactly like human taxi cab service!

Example

1. Central: Pick up passenger at 123 Garden St.?
2. Car A: Pick-up possible in 4 minutes!
3. Car B: Pick-up possible in 2 minutes!
4. Central: B is closer, car B will pick up.

Problem: Car better off lying about pick-up time $\rightarrow$ Passengers suffer!

Collaborative situation, but selfishness of agents (or their owners) interferes with system's optimality

This Lecture

1. How to act in presence of other (selfish) agents?
   - Maximize own utility/profit
   - Beat the opponent in competition
   - Must reason about opponents' behavior
2. Setting up multi-agent systems
   - Individuals' rationality $\not\rightarrow$ Group rationality
     Inherent problem in non-cooperative scenarios
   - Group rationality by changing individual's incentives

Strategic Games (Normal Form Games)

Players choose strategies, and get pay-offs accordingly:

		player 2	
		strategy C	strategy D
player 1	strategy A	p_{1AC}, p_{2AC}	$p_{1AD}, p_{2AD},$
	strategy B	$p_{1BC}, p_{2BC},$	p_{1BD}, p_{2BD}

(Examples are 2-player, but everything generalizes to any number.)

Dominated Strategies

Should Player 1 choose A or B?

	C	D
A	0,2	1,3
B	2,4	2,1

Strategy B is better for 1, no matter how 2 plays.

Assuming Player 1 is rational and will play B, Player 2 should play C (because $4 > 1$.)

Dominated Strategies

Definition

A strategy is strictly dominated if some other strategy is strictly better than it for every opponent strategy.

Theorem

If

1. *all players are rational,*
2. *all players know that all players are rational and that everybody know this (common knowledge), and*
3. *elimination of strictly dominated strategies leaves exactly one strategy for every agent,*

then these strategies will be played by the agents.

Prisoners' Dilemma

Two suspects are in detention, and asked to give information to police. Options are:

1. Cooperate with police, helping convict the other suspect on serious charges
2. Not cooperate, getting a small sentence

	shut up	cooperate
shut up	-1,-1	-3,0
cooperate	0,-3	-2,-2

Best if both suspects shut up, if commitment to this could be enforced.
But, assuming the other has already committed to a strategy, it is always best to cooperate!

Arms Race

Confrontation between two super-powers:

1. Arm: Acquire lots of weapons, exert threat
2. Don't arm

	don't arm	arm
don't arm	0,0	-3,1
arm	1,-3	-2,-2

Whatever the other does, it is always better to arm
However, (0,0) better than (-2,-2) ($\rightarrow$ Arms treaties)

Mixed Strategies

- When players have opposing goals, randomly selecting between strategies – which forces the opponent to randomize (and compromise his utility) – may produce the highest payoffs.
- example: Paper, Rock, Scissors
- example: bluffing and sand-bagging in Poker
  - With bad cards, bidding higher makes the opponent believe your cards are good
  - With good cards, not bidding higher makes the opponent believe your cards are bad

Both are part of optimal strategies for Poker

Examples of Mixed Strategies

Meeting

	A	B	C
A	1,1	0,0	0,0
B	0,0	1,1	0,0
C	0,0	0,0	1,1

versus...

Stalking

	A	B	C
A	-1,1	0,0	0,0
B	0,0	-1,1	0,0
C	0,0	0,0	-1,1

Paper, Rock, Scissors

	P	R	S
P	0,0	1,-1	-1,1
R	-1,1	0,0	1,-1
S	1,-1	-1,1	0,0

Formalization of Normal Form Games

Definition

1. There are players $1, \ldots, n$
2. Player i has a set of pure strategies A_i
3. A strategy profile is $(s_1, \ldots, s_n)$ such that $s_i \in A_i$ for every $i \in \{1, \ldots, n\}$
4. The utility $u_i(s_1, \ldots, s_n)$ of player i is expressed by a function $u_i : A_1 \times A_2 \times \cdots \times A_n \rightarrow \mathbb{R}$

Formalization of Normal Form Games

Definition

A mixed strategy σ is a probability distribution on the player's pure strategies A such that $\sum_{s \in A} \sigma(s) = 1$

Definition (Utility of Mixed Strategies)

The utility $u_i(\sigma_1, \ldots, \sigma_n)$ of player i w.r.t. a mixed strategy profile is

$$\sum_{(s_1,\ldots,s_n) \in A_1 \times \cdots \times A_n} u_i(s_1, \ldots, s_n)\sigma_1(s_1) \cdots \sigma_n(s_n)$$

Nash Equilibrium

Definition

A (mixed) strategy profile $(\sigma_1, \ldots, \sigma_n)$ is a Nash equilibrium if for all $i \in \{1, \ldots, n\}$,

$$u_i(\sigma_1, \ldots, \sigma_n) \geq u_i(\sigma_1, \ldots, \sigma_i', \ldots, \sigma_n) \text{ for all } \sigma_i' \neq \sigma_i$$

Nash Equilibrium: Intuitive meaning

NE is a (mixed) strategy profile in which no player can improve its pay-off by changing to some other strategy.

Nash Equilibrium

Disclaimer

It is not always clear how an NE could be reached and which NE would be played.
Nash equilibrium (without further assumptions) is realistic when

1. the agents cannot communicate,
2. the agents cannot bargain (re-distribute pay-offs),
3. game is played only once,
4. there is only one NE.

Nash Equilibrium: Properties

1. Every game has at least one NE
2. Not all games have pure-strategy NE
3. Every pure strategy in an NE is an equally good response to the opponents' strategies (randomization "only" needed to force opponents to randomize)

For 2-player zero-sum games, NE can be found with Linear Programming. More general games require more complex methods.

Example: Mixed Strategies

Bach or Stravinsky (aka Battle of the Sexes, with Ballet and Soccer)

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

Nash equilibria in pure strategies: $\langle B, B\rangle, \langle S, S\rangle$, with payoff $\langle 2, 1\rangle, \langle 1, 2\rangle$.
Nash equilibrium $\langle \sigma_1, \sigma_2\rangle$ in mixed strategies with

$$\begin{aligned} \sigma_1(B) &= \tfrac{2}{3} & \sigma_1(S) &= \tfrac{1}{3} \\ \sigma_2(B) &= \tfrac{1}{3} & \sigma_2(S) &= \tfrac{2}{3} \end{aligned}$$

and payoffs $\langle \frac{6}{9}, \frac{6}{9}\rangle$.