

CS 121

Multiple Agents

Lecture 11

Russell & Norvig ◊ Chapter 6, 17.6 - 17.8

Chris Archibald
July 29, 2009

Outline

- ▶ Examples
- ▶ Game theory
 - ▶ Definitions
 - ▶ Solution concepts: Equilibria
 - ▶ Interesting tidbits
 - ▶ Repeated games
- ▶ Social Choice
 - ▶ Definitions
 - ▶ Voting rules
 - ▶ Impossibilities
- ▶ Auctions
 - ▶ Different types
 - ▶ Equivalent types
- ▶ Other topics

Multiple Agents

Having multiple agents in the environment complicates things significantly. There are many different topics that arise in such settings:

- ▶ Decision making: How can our agent make *rational* decisions when the environment includes agents with opposing goals that will react to our decisions?
- ▶ Planning: How can we devise a plan that will take into account the fact that an agent might be trying to stop us? How can we plan the actions of a team of agents to accomplish some goal?
- ▶ Group decision making: How can our team of agents make some decision together when they all might have different views about the world and what is best?
- ▶ Interaction design: How should we design rules of interaction so that agents tell the truth and interact as we desire?
- ▶ Learning: How can I learn meaningful behaviour when who I am learning to play against is learning also?

Let's play some games

First game: Guess half the average

It will proceed as follows:

- ▶ 10 players each write down an integer between 0 and 100
- ▶ The average of the numbers is computed
- ▶ Whoever submitted the number closest to half the average wins the prize!

Let's play!

Half-the-average

- ▶ What do you think would happen if we played this game even longer?
- ▶ Would our strategies converge?
- ▶ How would you design an agent to play this game?
- ▶ What real-world situation could this apply to?

Game theory

Game theory is the mathematical study of agent interaction. Typically we assume the following:

- ▶ Each agent is rational (maximizes expected utility)
- ▶ Each agent knows that the other is rational
- ▶ All of this is common knowledge

Why are these assumptions made? Are they usually reasonable? Does this mean agents are selfish?

Aside: selfishness

- ▶ Sometimes the limited view of simply maximizing expected utility seems not to capture enough
- ▶ Answer: Modify the utility function to include anything else you want
- ▶ All of the complexity of modeling the situation is in the utility function
- ▶ Game theory typically take the utility function as a given

Normal-form games

The simplest type of game is called a **Normal-form game**. First a definition, then an example

A **Normal-form game** is defined by:

- ▶ A set of n players $I = \{1, \dots, n\}$.
- ▶ A set of actions. For each player i , their set of actions is A_i , with typical element a_i . We denote $A = \times_{i \in I} S_i$
- ▶ A payoff function. For each $i \in I$, a payoff function $u_i : S \mapsto \mathbb{R}$

Each player chooses an action, then they receive the payoff specified by their payoff function for that combination of actions

Strategies

We will generally call the manner in which a player chooses an action a **strategy**. There are two main types:

- ▶ Pure strategy: a player simply chooses an action
- ▶ Mixed strategy: a player randomly selects an action according to some distribution

It is easy to see that a pure strategy is a specific type of mixed strategy. We will see why this can be useful.

Normal-form game examples

		2	
		c	d
1	a	1 2 1 -1	
	b	-1 0 2 0	

Another example

		Jane	
		A	B
John	A	1 -1	-1 1
	B	-1 1	1 -1

The questions of game theory

Different areas of game theory address different questions:

- ▶ What should a rational agent do in this game?
- ▶ What are reasonable things for a rational agent to do in this game?
- ▶ What are things that a rational agent would never do in this game?
- ▶ What outcomes could be reached?
- ▶ Can we predict what a human would do in this game?
- ▶ Can we learn what we should do in this game?
- ▶ ...

What can we say about a game?

		2	
		c	d
1	a	1	2
	b	-1	0
		2	0

- ▶ We can compare different outcomes
- ▶ A **Pareto optimal** or **efficient** outcome is one in which no agent's utility can be improved without another agent's utility being decreased.
- ▶ But why would player's play this, especially if they only care about their own utility?

What can we say about a game?

		2	
		c	d
1	a	1	2
	b	-1	0
		2	0

Let's play this game and see what happens!

What can we say about a game?

		2	
		c	d
1	a	1, 1	2, -1
	b	-1, 2	0, 0

Notice, that no matter what the opponent does, player 1 is better off playing *a*. This is the easiest case, and is called a dominant strategy:

- ▶ A **dominant strategy** is a strategy that yields a higher expected utility than any other strategy, regardless of what the opponent is doing
- ▶ Can be weak or strong
- ▶ A **dominated strategy** is a strategy for which there exists another strategy which always does better

Game outcomes

- ▶ So, if each player has a dominant strategy, then a reasonable outcome would be the outcome where they each play this dominant strategy
- ▶ This is called a **dominant strategy equilibrium**. This is the easiest solution concept in game theory to justify and explain.
- ▶ So are we done?
- ▶ Sadly, not very many game situations have dominant strategies for the players. But, is it always the right answer?

The prisoner's dilemma

	Prisoner B Stays Silent	Prisoner B Betrays
Prisoner A Stays Silent	Each serves 6 months	Prisoner A: 10 years Prisoner B: goes free
Prisoner A Betrays	Prisoner A: goes free Prisoner B: 10 years	Each serves 5 years

The dominant strategy equilibrium is not Pareto efficient!

Best response

- ▶ What if an agent knew what the other agents were going to do in a game?
- ▶ Then he could select the strategy that would do best against the specific strategies of the opponents. This is called his **best-response**.
- ▶ A best response for player i is a function

$$BR_i : S_{-i} \mapsto S_i$$

, such that

$$u_i(BR_i(s_{-i}), s_{-i}) \geq u_i(s_i, s_{-i})$$

for all strategies $s_i \in S_i$.

Nash equilibrium

- ▶ Now we come to arguably the most famous solution concept of game theory – **Nash equilibrium**.
- ▶ A strategy profile s is a Nash equilibrium if for each player i , $s_i = BR_i(s_{-i})$

Intuition: If everyone knew what everyone else's strategy, they wouldn't change their own. A stable outcome of the game

Finding Nash equilibria

		Husband	
		LW	WL
Wife	LW	1 2	0 0
	WL	0 0	2 1
		Husband	

Half-the-average game

What is the Nash equilibrium for the Guess half-the-average game?

Nash equilibrium

- ▶ Named after John Nash, who proved in 1951 that a Nash equilibrium will exist in mixed strategies for any finite normal form game
- ▶ Won the Nobel Prize in Economics in 1994
- ▶ Was the subject of the movie “A Beautiful Mind”
- ▶ The concept of Nash equilibrium identifies a set of reasonable outcomes of a game

Coming to equilibrium

- ▶ Nash equilibrium (and other solution concepts) say something about reasonable results of the game, but nothing about how the players come to play it
- ▶ With a single play of a normal form game, there isn't really much in the way of "learning" to talk about
- ▶ What if we play a game over and over?

Repeated games

- ▶ Players compete in a normal form game, and then, play again
- ▶ Can be for some finite number of games or forever
- ▶ What kinds of strategies can players employ now?
- ▶ How should a rational player act?
- ▶ Does this change anything?

Repeated prisoner's dilemma

- ▶ Imagine that we repeat the prisoner's dilemma from before
- ▶ Can players ever cooperate?

Finite repetitions

- ▶ If we repeat the game a finite number of times, it is always best to defect
- ▶ We can see this by backwards induction
- ▶ In the last stage, the player's know that the game will end, and so it is best to defect.
- ▶ Then, in the second to last stage, knowing that both will defect in the last, it is best to defect . . .

What about infinite repetitions?

- ▶ Robert Axelrod studied this, by holding a tournament in early 1980's
- ▶ What program won?
- ▶ The simplest program entered. (4 lines of BASIC) ⇒ Tit-for-tat
- ▶ Cooperate first round, then do whatever the opponent did the previous round

Infinite repetition

- ▶ Ideally, we would like a good repeated-game-playing agent to learn how best to respond to its opponent
- ▶ There are many ways of doing this (fictitious play), but we won't go into them to much
- ▶ I will show one example of some difficulties that arise

Learning vs. teaching

		Driver 2	
		Straight	Swerve
Driver 1	Straight	0, 0	1, 3
	Swerve	3, 1	2, 2

- ▶ One player is an awesome learning agent, responding perfectly to the opponent
- ▶ The other player is crazy, irrational, and will drive straight every time
- ▶ What will happen?
- ▶ The crazy player wins!
- ▶ It taught the learning agent how to play so that it would win.
- ▶ Fine line between learning and teaching, and what can actually help an agent

Summary

Main ideas:

- ▶ Game theory is the mathematical study of interacting decision makers
- ▶ Basic definitions involved in a game
- ▶ Definition of Nash equilibrium
- ▶ How to find a Nash equilibrium in normal form games
- ▶ Feel for some of the questions that arise