

Algorithms 3: Game Theory

Loosely based on ECON 159 Open Yale Courses

Classes of interest

- CS4700: Artificial Intelligence
- ECON4020: Game Theory
- ORIE 4350: Game Theory

ECE 3400: Intelligent Physical Systems



Prisoner's Dilemma

- If you and your accomplice plead not guilty, you get a 2 year sentence
 - If you and your accomplice plead guilty, you get a 3 year sentence
 - If you plead guilty and your accomplice pleads not guilty, you get a 1 year sentence and your accomplice gets a 5 year sentence.
 - If your accomplice pleads guilty and you plead not guilty, you get a 5 year sentence and your accomplice get a 1 year sentence.
-
- *What would you do?*



"Well, if you ever want to talk, you know where to find me.
My door is never open."

End Game

5 weeks left till competition day

Coverage

The full mazes will be 9 x 9 squares. The robot that maps the most of the maze correctly (wrt to walls and gaps) in a given round will receive 15 points. All other robots will receive scaled values thereof.

Treasures

- For every treasure which is located correctly: 1 point
- For every treasure that is located and color-identified correctly: 1 point
- For every treasure that is located and shape-identified correctly: 1 point
- For every discovery of a treasure that is not there: -1 point
- The minimum score per round is 0 points; the maximum is 20 points.

How can you cheat the system?

- If no one moves!


BUT...

- If one robot moves, it gets 15+5 points and everyone else 0!

~~Prisoner's~~ Dilemma's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



		opponent	
		go	stay
me	go	5	20
	stay	0	15
		my points	


		opponent	
		go	stay
me	go	5	0
	stay	20	15
		opponent points	

		opponent		utilities
		go	stay	
me	go	5,5	20,0	
	stay	0,20	15,15	
		Superimposed "outcome matrix"		

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- If the opponent goes, and you go you get 5
- If the opponent goes, and you stay you get 0
- If the opponent stays, and you go you get 20
- If the opponent stays, and you stay you get 15
- *Going always yields the bigger payoff!*

		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- **Lesson:**
- The go-strategy strictly dominates the stay-strategy, if the pay off from going is strictly greater than from staying, regardless of what the opponent does.
- 1. Never play a strictly dominated strategy
 - *(Play the strictly dominating strategy)*

		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- *Is this really true?*
- ...if we all go, we all get lower grades
- Your reasoning does not affect your opponent
 - (If you could affect the opponents choice, it would make more sense to have them stay so that you could go.)

		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



• **Lessons:**

1. Never play a strictly dominated strategy
2. Dominant strategies can lead to inefficient outcomes

		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma


- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- *What if we all just agree to stay?*
 - Without reinforcement, how can you trust that everyone will actually stay?

Arduino Schwarzenegger
(Rage Against the Machines)



		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma


- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- *What if we all just agree to stay?*
 - Without reinforcement, how can you trust that everyone will actually stay?
- *What if the pay-offs look*

Gary the robot
(The Good Noodles)



		opponent	
		go	stay
me	go	5,5	20,0
	stay	0,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- If I choose to go, and everyone else stays I feel guilty ($20 - 25 = -5$)
- If I choose to stay, and everyone else goes, I feel indignant ($0 - 15 = -15$)

Which strategy is dominating?

- Neither
- **Lesson:**
 - Payoffs matter

		opponent	
		go	stay
me	go	5,5	-5,-15
	stay	-15,-5	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



What if you are on Rage Against the Machines and playing the Good Noodles?

		opponent	
		go	stay
me	go	5,5	-5,-15
	stay	-15,-5	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



What if you are on Rage Against the Machines and playing the Good Noodles?

- Go always gives a higher payoff
- Go strictly dominates staying

The Good Noodles

		go	stay
Rage Against the Machines	go	5,5	20,-15
	stay	0,-5	15,15


utilities

Superimposed "outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- *What if you are on Rage Against the Machines and playing the Good Noodles?*
 - Go always gives a higher payoff
 - Go strictly dominates staying
- *What if you are a Good Noodle and playing Rage Against the Machines?*

The Good Noodles

	go	stay
Rage Against the Machines go	5,5	20,-15
stay	0,-5	15,15


utilities

Superimposed "outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



- *What if you are on Rage Against the Machines and playing the Good Noodles?*
 - Go always gives a higher payoff
 - Go strictly dominates staying
- *What if you are a Good Noodle and playing Rage Against the Machines?*
 - You know your opponent will go, so your best option is to go!

Rage Against the Machines

		Rage Against the Machines	
		go	stay
The Good Noodles	go	5,5	-5,0
	stay	-15,20	15,15


utilities

Superimposed
"outcome matrix"

~~Prisoner's~~ Student's Dilemma

- If you both stay, both get 15 points
- If you both go, both get >0 points (5)
- If you go and your opponents stay, you get 20 points and they get 0
- If you stay and your opponents go, you get 0 and your opponents get 20 points

What do you do?



• *Lessons:*

1. Never play a strictly dominated strategy
2. Dominant strategies can lead to inefficient outcomes
3. Decide based on what you think the other robot will do
4. Never play a weakly dominated strategy

Rage Against the Machines

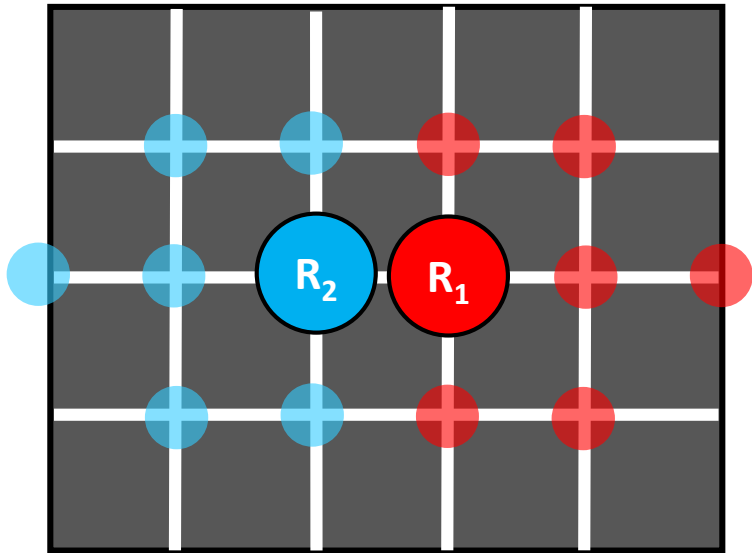

		go	stay
The Good Noodles	go	5,5	-5,0
	stay	-15,20	15,15

utilities

Superimposed
"outcome matrix"

Similar Situations in the Competition?

- A robot encountering another robot
- *Fill in the utilities of the outcome matrix, and decide which is the better choice*



		R_2	
		Reverse	Stay
R_1	Reverse	-1,-1	-1,1
	Stay	1,-1	0,0

We should both stay!

Prisoner's Dilemma

- If you and your accomplice plead not guilty, you get a 2 year sentence
- If you and your accomplice plead guilty, you get a 3 year sentence
- If you plead guilty and your accomplice pleads not guilty, you get a 1 year sentence and your accomplice gets a 5 year sentence.
- If your accomplice pleads guilty and you plead not guilty, you get a 5 year sentence and your accomplice get a 1 year sentence.

- **Real World Answers**

- 70% choose to plead *guilty*
- 30% choose to plead *not guilty* (dominated strategy)

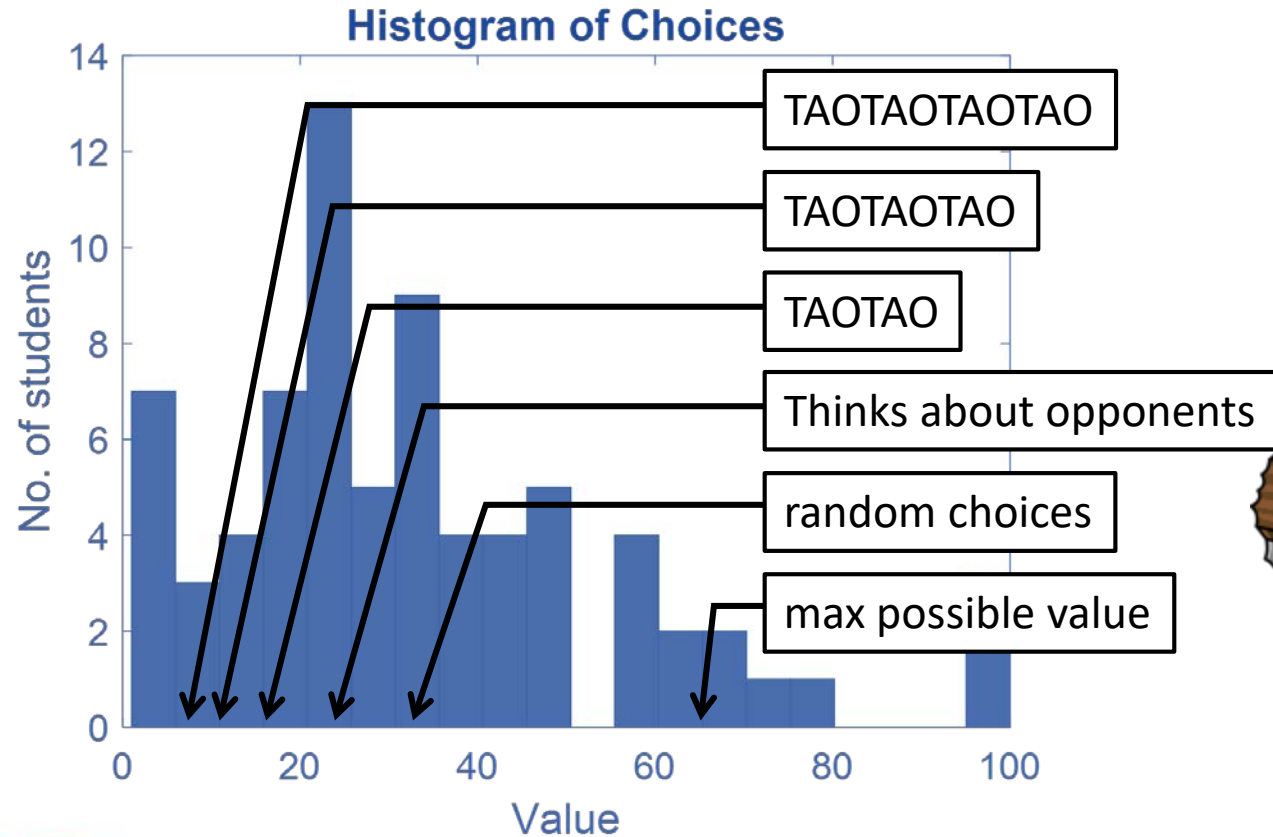
- **Why?**

- Friendship
- Ignorant
- Considering long term gain

		Henry	
		Not Guilty	Guilty
Dave	Not Guilty	2 Years	5 Years, 1 Yr.
	Guilty	5 Years, 1 Yr.	3 Years

Game Theory

- Pick a whole number between 1 and 100.
- The winner is the person who picks the value which is closest to two thirds of the class average.
- Class average: 33.5
- $2/3s$: 22.34
- Winners:
 - Alexander Hatzis
 - Gauri Jain
 - Tyler Sherman



Go Build Robots!

