

CS378

Autonomous Multiagent Systems

Spring 2005

Prof: Peter Stone
TA: Nate Kohl

Department of Computer Sciences
The University of Texas at Austin

Week 13b: Thursday, April 20th

Good Afternoon, Colleagues

Are there any questions?

Class Discussion

Mickey Ristroph on human modeling

Reinforcement Learning

- Slides from Tom Mitchell's ML book

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

- Nash equilibrium?

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

- Nash equilibrium?
- Action 2 is dominant for Player 1. End of story?

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

- Nash equilibrium?
- Action 2 is dominant for Player 1. End of story?
- What would you do as player 2?

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

- Nash equilibrium?
- Action 2 is dominant for Player 1. End of story?
- What would you do as player 2?
- **Threats** can stabilize a non-equilibrium strategy
- Change the **best response** of the other agent

Stackelburg Game

		Player 2	
		Action 1	Action 2
Player 1	Action 1	1,0	3,2
	Action 2	2,1	4,0

- Nash equilibrium?
- Action 2 is dominant for Player 1. End of story?
- What would you do as player 2?
- **Threats** can stabilize a non-equilibrium strategy
- Change the **best response** of the other agent

Threats slides

Discussion

- How useful is the concept of Nash equilibrium?

Discussion

- How useful is the concept of Nash equilibrium?
- Shoham:
 - 0-sum = single agent problem
 - common payoff = search for pareto optimum

Discussion

- How useful is the concept of Nash equilibrium?
- Shoham:
 - 0-sum = single agent problem
 - common payoff = search for pareto optimum
 - General sum is the interesting case:

Discussion

- How useful is the concept of Nash equilibrium?
- Shoham:
 - 0-sum = single agent problem
 - common payoff = search for pareto optimum
 - General sum is the interesting case:
 - Learning in an environment with other, unknown, independent agents who may also be learning

Discussion

- How useful is the concept of Nash equilibrium?
- Shoham:
 - 0-sum = single agent problem
 - common payoff = search for pareto optimum
 - General sum is the interesting case:
 - Learning in an environment with other, unknown, independent agents who may also be learning
 - Need to do well against some set of agents, never too poorly, and well against yourself.