CS344M
Autonomous Multiagent Systems

Prof: Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Are there any questions?

Logistics

- All readings up

Logistics

- All readings up
  - Next Tuesday: guest lecture

Logistics

- All readings up
  - Next Tuesday: guest lecture
  - Responses also to Shivaram Kalyanakrishnan

Logistics

- All readings up
  - Next Tuesday: guest lecture
  - Responses also to Shivaram Kalyanakrishnan
- More reflections on peer reviews?

Logistics

- All readings up
  - Next Tuesday: guest lecture
  - Responses also to Shivaram Kalyanakrishnan
- More reflections on peer reviews?
- Final projects due in 2 weeks!

Reading Overview — Vidal and Durfee

Recursive Modeling Method

- What should I do?

Reading Overview — Vidal and Durfee

Recursive Modeling Method

- What should I do?
- What should I do given what I think you'll do?

Reading Overview — Vidal and Durfee

Recursive Modeling Method

- What should I do?
- What should I do given what I think you'll do?
- What should I think you'll do given what I think you think I'll do?

Reading Overview — Vidal and Durfee

Recursive Modeling Method

- What should I do?
- What should I do given what I think you'll do?
- What should I think you'll do given what I think you think I'll do?
- etc.

Prediction Method

- Rely on communication

Prediction Method

- Rely on communication
  - What to say? What to trust?

Prediction Method

- Rely on communication
  - What to say? What to trust?
- Watch for patterns of others

Prediction Method

- Rely on communication
  - What to say? What to trust?
- Watch for patterns of others
  - Might have incorrect expectations, especially if environment changes

Prediction Method

- Rely on communication
  - What to say? What to trust?
- Watch for patterns of others
  - Might have incorrect expectations, especially if environment changes
- Use deeper models
  - Includes physical *and* mental states

Prediction Method

- Rely on communication
  - What to say? What to trust?
- Watch for patterns of others
  - Might have incorrect expectations, especially if environment changes
- Use deeper models
  - Includes physical *and* mental states
  - Could be computationally expensive

Types of models

Example: pursuit task

No-information: Random choice

Types of models

Example: pursuit task

No-information: Random choice

Sub-intentional: Not rational

Types of models

Example: pursuit task

No-information: Random choice

Sub-intentional: Not rational

Intentional: Others use same model

Lessons

- Modeling can help
- There is a lot of useless information in recursive models
- Approximations (limited rationality) can be useful

Tracking Dynamic Team Activity

- Use your own plans to model others

Tracking Dynamic Team Activity

- Use your own plans to model others
- Use explicit team operators

Tracking Dynamic Team Activity

- Use your own plans to model others
- Use explicit team operators
  - Introduces challenges of role assignments, and
  - Minimum cost repair

Tracking Dynamic Team Activity

- Use your own plans to model others
- Use explicit team operators
  - Introduces challenges of role assignments, and
  - Minimum cost repair
- Assume agent is using a plan that you could use,
  - But not modeling you

Tracking Dynamic Team Activity

- Use your own plans to model others
- Use explicit team operators
  - Introduces challenges of role assignments, and
  - Minimum cost repair
- Assume agent is using a plan that you could use,
  - But not modeling you
- Act based on assumed actions of others

Other papers you read

- Dustin: "Opponent Behavior Recognition for Real-Time Strategy Games"

Other papers you read

- Dustin: "Opponent Behavior Recognition for Real-Time Strategy Games"
- Aaron: "Distinguishing Between Intentional and Unintentional Sequences of Actions"

Where do Models Come From

Observation:

- Tambe and RMM: use existing model
  - No building a model

Where do Models Come From

Observation:

- Tambe and RMM: use existing model
  - No building a model

What if we can't build a full model in advance?

Where do Models Come From

Observation:

- Tambe and RMM: use existing model
  - No building a model

What if we can't build a full model in advance?

- What are some incremental approaches for building a predictive model?

Play me at RoShamBo

- Rock beats scissors
- Scissors beats paper
- Paper beats rock

Play me at RoShamBo

- Rock beats scissors
- Scissors beats paper
- Paper beats rock

- What is your strategy before modeling me?

Play me at RoShamBo

- Rock beats scissors
- Scissors beats paper
- Paper beats rock

- What is your strategy before modeling me?

- What is your strategy after modeling me?

Play me at RoShamBo

- Rock beats scissors
- Scissors beats paper
- Paper beats rock

- What is your strategy before modeling me?
- What is your strategy after modeling me?
- Am I modeling you?

Play me at RoShamBo

- Rock beats scissors
- Scissors beats paper
- Paper beats rock

- What is your strategy before modeling me?
- What is your strategy after modeling me?
- Am I modeling you?
- Would your end strategy change if I can?

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