CS344M
Autonomous Multiagent Systems

Patrick MacAlpine

Department or Computer Science
The University of Texas at Austin

Good Afternoon, Colleagues

Are there any questions?

Logistics

- Assignment 4 due today

Logistics

- Assignment 4 due today
- FAI talk Monday in GDC 6.302, 11am

Logistics

- Assignment 4 due today

- FAI talk Monday in GDC 6.302, 11am
  - Anca Dragan: “Towards Robotics Algorithms that Reason about People”

Logistics

- Assignment 4 due today
- FAI talk Monday in GDC 6.302, 11am
  - Anca Dragan: "Towards Robotics Algorithms that Reason about People"
- Final project proposal assigned

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Progress Report (11/5): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Progress Report (11/5): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Peer Review (11/12): review 2 progress reports

- Clear? Suggestions?; graded on writing and feedback quality

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Progress Report (11/5): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Peer Review (11/12): review 2 progress reports

- Clear? Suggestions?; graded on writing and feedback quality

Team (12/1): source + binaries

- The tournament entry; make sure it runs!

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Progress Report (11/5): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Peer Review (11/12): review 2 progress reports

- Clear? Suggestions?; graded on writing and feedback quality

Team (12/1): source + binaries

- The tournament entry; make sure it runs!

Final Report (12/3): 8+ pages

- A term paper; the main component of your grade

Final Projects

Proposal (10/8): 3+ pages

- What you're going to do; graded on writing

Progress Report (11/5): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Peer Review (11/12): review 2 progress reports

- Clear? Suggestions?; graded on writing and feedback quality

Team (12/1): source + binaries

- The tournament entry; make sure it runs!

Final Report (12/3): 8+ pages

- A term paper; the main component of your grade

Tournament (12/9 7-10PM): nothing due

- Oral presentation

Final Project info

- All writing is individual!

Final Project info

- All writing is individual!
- Two hard copies and one electronic copy

Final Project info

- All writing is individual!
- Two hard copies and one electronic copy
- Due at beginning of class

Final Project info

- All writing is individual!
- Two hard copies and one electronic copy
- Due at beginning of class
- One idea: Re-implement an idea from one of the readings

Final Project info

- All writing is individual!
- Two hard copies and one electronic copy
- Due at beginning of class
- One idea: Re-implement an idea from one of the readings
- Be careful with machine learning

Final Project info

- All writing is individual!
- Two hard copies and one electronic copy
- Due at beginning of class
- One idea: Re-implement an idea from one of the readings
- Be careful with machine learning
- Example final report on website

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley et al: Coach competition, extracting models

Kuhlmann et al: Learning for coaching

Wihthopf and Reidmiller: Reinforcement learning

MacAlpine, Price, and Stone: Role assignment

MacAlpine, Depinet, and Stone: Overlapping layered learning

Architecture for Action Selection

- leading passes

Architecture for Action Selection

- leading passes
- force fields for movement

Architecture for Action Selection

- leading passes
- force fields for movement
- (other slides, video)

Architecture for Action Selection

- leading passes
- force fields for movement
- (other slides, video)
- downsides

Architecture for Action Selection

- leading passes
- force fields for movement
- (other slides, video)
- downsides
- Keepaway

Architecture for Action Selection

- leading passes
- force fields for movement
- (other slides, video)
- downsides
- Keepaway
- (3D slides on action selection)

Coaching

- Learn best strategy to play a fixed team

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations
- Learn when successful teams passed/kicked

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations
- Learn when successful teams passed/kicked
- Learn when opponent will pass and try to block

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations
- Learn when successful teams passed/kicked
- Learn when opponent will pass and try to block
- What if players switch roles?

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations
- Learn when successful teams passed/kicked
- Learn when opponent will pass and try to block
- What if players switch roles?
- Why just imitate another team?

Coaching

- Learn best strategy to play a fixed team
- Give high level advice to players at low frequency
- Focus on learning formations
- Learn when successful teams passed/kicked
- Learn when opponent will pass and try to block
- What if players switch roles?
- Why just imitate another team?
- Other slides

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)
- SCRAM: Scaleable Collision-avoiding Role Assignment with Minimal-makespan

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)
- SCRAM: Scaleable Collision-avoiding Role Assignment with Minimal-makespan
  - Scaleable (computable in polynomial time)

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)
- SCRAM: Scaleable Collision-avoiding Role Assignment with Minimal-makespan
  - Scaleable (computable in polynomial time)
  - Avoids collisions between agents

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)
- SCRAM: Scaleable Collision-avoiding Role Assignment with Minimal-makespan
  - Scaleable (computable in polynomial time)
  - Avoids collisions between agents
  - Minimizes the makespan (agents reach all target positions in as little time as possible)

Role Assignment

- Find 1 to 1 mapping of agents to roles (target positions)
- SCRAM: Scaleable Collision-avoiding Role Assignment with Minimal-makespan
  - Scaleable (computable in polynomial time)
  - Avoids collisions between agents
  - Minimizes the makespan (agents reach all target positions in as little time as possible)
- (slides)

Reinforcement Learning

- Extend to grid soccer

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options
- Successfully learn the task, use for some of team behavior

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options
- Successfully learn the task, use for some of team behavior
- However, takes 12 million actions to learn

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options
- Successfully learn the task, use for some of team behavior
- However, takes 12 million actions to learn
- Current state of the art: MAXQ-OP

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options
- Successfully learn the task, use for some of team behavior
- However, takes 12 million actions to learn
- Current state of the art: MAXQ-OP
  - "Towards a Principled Solution to Simulated Robot Soccer" - Aijun Bai, Feng Wu, and Xiaoping Chen

Reinforcement Learning

- Extend to grid soccer
- Large state space, joint actions
- Address this with state aliasing, options
- Successfully learn the task, use for some of team behavior
- However, takes 12 million actions to learn
- Current state of the art: MAXQ-OP
  - "Towards a Principled Solution to Simulated Robot Soccer" - Aijun Bai, Feng Wu, and Xiaoping Chen
  - Used by 2D champion WrightEagle team