

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

- How to read a research paper

Logistics

- How to read a research paper
 - Some have too few details...

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted
- FAI talk Thursday in ACES 2.402, 2pm

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted
- FAI talk Thursday in ACES 2.402, 2pm
 - Philip Resnick: "Translation as Collaboration"

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted
- FAI talk Thursday in ACES 2.402, 2pm
 - Philip Resnick: "Translation as Collaboration"
- Thesis proposal Wednesday in ACES 3.116, 8:45am

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted
- FAI talk Thursday in ACES 2.402, 2pm
 - Philip Resnick: "Translation as Collaboration"
- Thesis proposal Wednesday in ACES 3.116, 8:45am
 - Todd Hester: "Temporal Difference Reinforcement Learning in Time-Constrained Domains"

Logistics

- How to read a research paper
 - Some have too few details...
 - Others have too many.
- Next week's readings posted
- FAI talk Thursday in ACES 2.402, 2pm
 - Philip Resnick: "Translation as Collaboration"
- Thesis proposal Wednesday in ACES 3.116, 8:45am
 - Todd Hester: "Temporal Difference Reinforcement Learning in Time-Constrained Domains"
- Final project proposal assigned

Final Projects

Proposal (10/7): 3+ pages

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

Final Projects

Proposal (10/7): 3+ pages

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

Progress Report (11/4): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Final Projects

Proposal (10/7): 3+ pages

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

Progress Report (11/4): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Team (11/30): source + binaries

- The tournament entry; make sure it runs!

Final Projects

Proposal (10/7): 3+ pages

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

Progress Report (11/4): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Team (11/30): source + binaries

- The tournament entry; make sure it runs!

Final Report (12/2): 8+ pages

- A term paper; the main component of your grade

Final Projects

Proposal (10/7): 3+ pages

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

Progress Report (11/4): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Team (11/30): source + binaries

- The tournament entry; make sure it runs!

Final Report (12/2): 8+ pages

- A term paper; the main component of your grade

Tournament (TBA): nothing due

- Oral presentation

Final Projects

Proposal (10/7): 3+ pages

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

Progress Report (11/4): 5+ pages + binaries + logs

- What you've been doing; graded on writing

Team (11/30): source + binaries

- The tournament entry; make sure it runs!

Final Report (12/2): 8+ pages

- A term paper; the main component of your grade

Tournament (TBA): nothing due

- Oral presentation

Due at beginning of classes

Overview of the Readings

Darwin: genetic programming approach

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley: Coach competition, extracting models

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley: Coach competition, extracting models

Kuhlmann: Learning for coaching

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley: Coach competition, extracting models

Kuhlmann: Learning for coaching

Riedmiller05: Reinforcement learning

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley: Coach competition, extracting models

Kuhlmann: Learning for coaching

Riedmiller05: Reinforcement learning

Graf: Closed loop walk on real robots

Overview of the Readings

Darwin: genetic programming approach

Stone and McAllester: Architecture for action selection

Riley: Coach competition, extracting models

Kuhlmann: Learning for coaching

Riedmiller05: Reinforcement learning

Graf: Closed loop walk on real robots

Shafii: Open loop walk in 3d simulator

Evolutionary Computation

- Motivated by biological evolution: GA, GP

Evolutionary Computation

- Motivated by biological evolution: GA, GP
- Search through a space

Evolutionary Computation

- Motivated by biological evolution: GA, GP
- Search through a space
 - Need a **representation, fitness function**
 - Probabilistically apply search operators to set of points in search space

Evolutionary Computation

- Motivated by biological evolution: GA, GP
- Search through a space
 - Need a **representation, fitness function**
 - Probabilistically apply search operators to set of points in search space
- Randomized, parallel hill-climbing through space

Evolutionary Computation

- Motivated by biological evolution: GA, GP
- Search through a space
 - Need a **representation, fitness function**
 - Probabilistically apply search operators to set of points in search space
- Randomized, parallel hill-climbing through space
- Learning is an optimization problem (fitness)

Evolutionary Computation

- Motivated by biological evolution: GA, GP
- Search through a space
 - Need a **representation, fitness function**
 - Probabilistically apply search operators to set of points in search space
- Randomized, parallel hill-climbing through space
- Learning is an optimization problem (fitness)

Some slides from *Machine Learning* (Mitchell, 1997)

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction
- Evolves whole teams — lexicographic fitness function

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction
- Evolves whole teams — lexicographic fitness function
- Evolved on huge (at the time) hypercube

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction
- Evolves whole teams — lexicographic fitness function
- Evolved on huge (at the time) hypercube
- Lots of spinning, but figured out dribbling, offsides

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction
- Evolves whole teams — lexicographic fitness function
- Evolved on huge (at the time) hypercube
- Lots of spinning, but figured out dribbling, offsides
- 1-1-1 record. Tied a good team, but didn't advance

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by my detailed team construction
- Evolves whole teams — lexicographic fitness function
- Evolved on huge (at the time) hypercube
- Lots of spinning, but figured out dribbling, offsides
- 1-1-1 record. Tied a good team, but didn't advance
- Success of the method, but not pursued

Architecture for Action Selection

- (other slides, video)

Architecture for Action Selection

- (other slides, video)
- downsides

Architecture for Action Selection

- (other slides, video)
- downsides
- don't call "action selection"