

CS344M

Autonomous Multiagent Systems

Patrick MacAlpine

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

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 et al: Coach competition, extracting models

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

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

Wihthop and Reidmiller: Reinforcement learning

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

Wihthop and Reidmiller: Reinforcement learning

MacAlpine, Price, and Stone: Role assignment

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

Wihthop and Reidmiller: Reinforcement learning

MacAlpine, Price, and Stone: Role assignment

MacAlpine, Depinet, and Stone: Overlapping layered learning

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 Peter's detailed team construction

Darwin United

- More ambitious follow-up to Luke, 97 (made 2nd round)
- Motivated in part by Peter's 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 Peter's detailed team construction
- Evolves whole teams — lexicographic fitness function
- Lots of spinning, but figured out dribbling, offsides

Darwin United

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

Overlapping Layered Learning

- Machine learning paradigms (not algorithms)

Overlapping Layered Learning

- Machine learning paradigms (not algorithms)
- Useful for learning complex skills that *work well together*

Overlapping Layered Learning

- Machine learning paradigms (not algorithms)
- Useful for learning complex skills that *work well together*
- (slides)