

CS394R
Reinforcement Learning:
Theory and Practice
Fall 2007

Peter Stone

Department of Computer Sciences
The University of Texas at Austin

Good Afternoon Colleagues

- Are there any questions?

Logistics

- Final project topics

Logistics

- Final project topics
- Papers on web page

The Week 0 Task?

The Week 0 Task?

- Continuous state, noisy actions

The Week 0 Task?

- Continuous state, noisy actions
- What's needed to do the normal thing?

The Week 0 Task?

- Continuous state, noisy actions
- What's needed to do the normal thing?
- How can we get it?

The Week 0 Task?

- Continuous state, noisy actions
- What's needed to do the normal thing?
- How can we get it?
- Keepaway example

Class discussion – Decision trees for FA

- Decision Trees
- T-test

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?
- Always decrease alpha?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?
- Always decrease alpha?
- Continuous actions?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?
- Always decrease alpha?
- Continuous actions?
- More info about bootstrapping successes?
- bias in (8.3)?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?
- Always decrease alpha?
- Continuous actions?
- More info about bootstrapping successes?
- bias in (8.3)?
- Baird - why on-policy converges
- Counterexamples rare or common?

Class discussion – Decision trees for FA

- Why these encodings (over linear)? - why not just NN?
- Mountain Car - is it worth going to 810 features?
- What makes an FA good/bad for RL?
- Ensembles for generalizing?
- Can you extract rules for generalization?
- Always decrease alpha?
- Continuous actions?
- More info about bootstrapping successes?
- bias in (8.3)?
- Baird - why on-policy converges
- Counterexamples rare or common?

Class discussion – Kinerva coding

- "Receptive field"