

CS395T
Reinforcement Learning:
Theory and Practice
Fall 2004

Peter Stone

Department of Computer Sciences
The University of Texas at Austin

Week10a: Tuesday, November 2nd

Good Afternoon Fellow Constituents

- Are there any questions?

Logistics

- Readings all posted

Logistics

- Readings all posted
 - Slightly different information types
 - Sparser response reactions

Logistics

- Readings all posted
 - Slightly different information types
 - Sparser response reactions
- Use as input for your final projects

Chapter 10 discussion

- p. 259 bottom: preparing to do other things?

Chapter 10 discussion

- p. 259 bottom: preparing to do other things?
- Relational MDPs

Chapter 11 discussion

- Self-play vs. playing vs. experts

Chapter 11 discussion

- Self-play vs. playing vs. experts
- Samuel's backup

Chapter 11 discussion

- Self-play vs. playing vs. experts
- Samuel's backup
- Why multiple tilings on the acrobat task?

Chapter 11 discussion

- Self-play vs. playing vs. experts
- Samuel's backup
- Why multiple tilings on the acrobat task?
- Channel allocation beats everything?

Chapter 11 discussion

- Self-play vs. playing vs. experts
- Samuel's backup
- Why multiple tilings on the acrobat task?
- Channel allocation beats everything?
- Planning: action sampling, TD backwards (286)

Chapter 11 discussion

- Self-play vs. playing vs. experts
- Samuel's backup
- Why multiple tilings on the acrobat task?
- Channel allocation beats everything?
- Planning: action sampling, TD backwards (286)
- How come so much NN FA in the chapter?