

CS395T
Reinforcement Learning:
Theory and Practice
Fall 2004

Peter Stone

Department of Computer Sciences
The University of Texas at Austin

Week6a: Tuesday, October 5th

Good Afternoon Colleagues

- Are there any questions?

Logistics

- Final project topics

Logistics

- Final project topics
- The TBA weeks

Logistics

- Final project topics
- The TBA weeks
- Programming projects public?

Survey Reactions

- Thanks!

Survey Reactions

- Thanks!
- Discussion moderation
- Focus on week 0 problem

Survey Reactions

- Thanks!
- Discussion moderation
- Focus on week 0 problem
- Email follow-ups of things not finished in class

Survey Reactions

- Thanks!
- Discussion moderation
- Focus on week 0 problem
- Email follow-ups of things not finished in class
- More alternatives to the book

Discussion Points

- N-step return in week 0 task (on-line)

Discussion Points

- N-step return in week 0 task (on-line)
- On-line vs. off-line in week 0 task

Discussion Points

- N-step return in week 0 task (on-line)
- On-line vs. off-line in week 0 task
- $TD(\lambda)$ on week 0 task

Discussion Points

- N-step return in week 0 task (on-line)
- On-line vs. off-line in week 0 task
- $TD(\lambda)$ on week 0 task
- Intuition for equation (7.2)

Discussion Points

- N-step return in week 0 task (on-line)
- On-line vs. off-line in week 0 task
- $TD(\lambda)$ on week 0 task
- Intuition for equation (7.2)
- Exercises 7.1–7.3

Class Discussion

- Sit on N-step returns vs. $TD(\lambda)$