CS395T
Reinforcement Learning: Theory and Practice
Fall 2004

Peter Stone

Department or Computer Sciences
The University of Texas at Austin

Week11a: Tuesday, November 9th

Good Afternoon Colleagues

- Are there any questions?

Options

- Extension of RL to temporal abstraction

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...
  - ...Week 0 task!

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...
  - ...Week 0 task!
  - p. 14?

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...
  - ...Week 0 task!
  - p. 14?
- They don't address **what** temporal abstraction to use — they just show how it can fit into the RL formalism

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...
  - ...Week 0 task!
  - p. 14?
- They don't address **what** temporal abstraction to use — they just show how it can fit into the RL formalism
  - Why couldn't it before?

Options

- Extension of RL to temporal abstraction
- State abstraction vs. temporal abstraction...
  - ...Week 0 task!
  - p. 14?
- They don't address **what** temporal abstraction to use — they just show how it can fit into the RL formalism
  - Why couldn't it before?
- Markov vs. Semi-markov:
  - states, actions
  - mapping from (s, a) to expected discounted reward
  - well-defined distribution of next state, transit time

Student-led Discussion

- Alex on discovering options

Discussion Points

- Why cliffs on p. 23 ?

Discussion Points

- Why cliffs on p. 23 ?
- Why lines cross on p. 27?

Discussion Points

- Why cliffs on p. 23 ?
- Why lines cross on p. 27?
- Does it ever make sense to have a probabilistic beta?
- What about non-terminating options?

Discussion Points

- Why cliffs on p. 23 ?
- Why lines cross on p. 27?
- Does it ever make sense to have a probabilistic beta?
- What about non-terminating options?
- Conclusion: "If the temporally extended actions are modeled as options, then perhaps the models of the options correspond well to these perceptions."