

CS394R

Reinforcement Learning: Theory and Practice

Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Afternoon Colleagues

- Are there any questions?

Logistics

- Next 2 weeks of readings are up

Logistics

- Next 2 weeks of readings are up
- Final project due in 3 1/2 weeks!

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

Discussion Points

- Are composed options *always* semi-Markov?

Discussion Points

- Are composed options *always* semi-Markov?
- What happens when initial value functions are optimistic?
(slides)

Discussion Points

- Are composed options *always* semi-Markov?
- What happens when initial value functions are optimistic?
(slides)
- Option discovery (class discussion)

Discussion Points

- Are composed options *always* semi-Markov?
- What happens when initial value functions are optimistic? (slides)
- Option discovery (class discussion)
 - bottleneck states
 - novelty
 - changed useful state abstractions (slides)

Discussion Points

- Are composed options *always* semi-Markov?
- What happens when initial value functions are optimistic? (slides)
- Option discovery (class discussion)
 - bottleneck states
 - novelty
 - changed useful state abstractions (slides)
- Options with function approximation possible?