

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 projects due on Thursday at 12:30
- Hard and soft copies

Overview

- Requires a method for updating a function approximator with knowledge (advice)

Overview

- Requires a method for updating a function approximator with knowledge (advice)
- KBANN already existed (but needed some extensions)

Overview

- Requires a method for updating a function approximator with knowledge (advice)
- KBANN already existed (but needed some extensions)
- In principle, other FAs could be used too

Some questions

- How are advice weights set? (Why does it matter?)

Some questions

- How are advice weights set? (Why does it matter?)
(Fig. 4, footnote 1, sec. 5)

Some questions

- How are advice weights set? (Why does it matter?)
(Fig. 4, footnote 1, sec. 5)
- Doesn't growing the network slow down learning?

Some questions

- How are advice weights set? (Why does it matter?)
(Fig. 4, footnote 1, sec. 5)
- Doesn't growing the network slow down learning?
- Freezing the network methodology - still on-line?

Some questions

- How are advice weights set? (Why does it matter?)
(Fig. 4, footnote 1, sec. 5)
- Doesn't growing the network slow down learning?
- Freezing the network methodology - still on-line?
- Would forgetting advice be as easy as they make it sound?