

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

- Reminder: more readings due on Thursday
 - Responses due 9pm Wed.
- Make progress on final projects!!

Background

- Backgammon
- NNs

How does it work?

- What's the role of the NN?

How does it work?

- What's the role of the NN?
- What's the input representation?

How does it work?

- What's the role of the NN?
- What's the input representation?
 - Truncated unary encoding (units for 1,2,3,4+ pieces on each space)
 - eventually: strength of blockade, probability of being hit

How does it work?

- What's the role of the NN?
- What's the input representation?
 - Truncated unary encoding (units for 1,2,3,4+ pieces on each space)
 - eventually: strength of blockade, probability of being hit
- Is it Sarsa or Q-learning?

How does it work?

- What's the role of the NN?
- What's the input representation?
 - Truncated unary encoding (units for 1,2,3,4+ pieces on each space)
 - eventually: strength of blockade, probability of being hit
- Is it Sarsa or Q-learning?
- What's lambda?

What does the NN learn?

An examination of the input-to-hidden weights in this network revealed interesting spatially organized patterns of positive and negative weights, roughly corresponding to what a knowledge engineer might call useful features for game play.

What does the NN learn?

An examination of the input-to-hidden weights in this network revealed interesting spatially organized patterns of positive and negative weights, roughly corresponding to what a knowledge engineer might call useful features for game play.

- Learn non-linear concepts?

Class Discussion - Anand

- What's the role of coevolution?