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 week's readings: abstraction

Logistics

- Next week's readings: abstraction
- Make progress on final projects!!

Background

- Bayes Nets, DBNs, CPTs (ppt)

Background

- Bayes Nets, DBNs, CPTs (ppt)
- Q-learning vs. RMax (videos)

Background

- Bayes Nets, DBNs, CPTs (ppt)
- Q-learning vs. RMax (videos)
  - Rmax (and SLF-Rmax) not built to be practical

Background

- Bayes Nets, DBNs, CPTs (ppt)
- Q-learning vs. RMax (videos)
  - Rmax (and SLF-Rmax) not built to be practical
- Factored state space (pdf)
- Structure learning

Background

- Bayes Nets, DBNs, CPTs (ppt)
- Q-learning vs. RMax (videos)
  - Rmax (and SLF-Rmax) not built to be practical
- Factored state space (pdf)
- Structure learning
  - Vs. value function approximation

Basic Structure-Learning Algoirthm

- Just one bit to predict (single factor, single action)

Basic Structure-Learning Algoirthm

- Just one bit to predict (single factor, single action)
- Just dependent on one bit
- Independent factors (no synchronic links)

Discussion

- What's more interesting? Theoretically grounded algorithms? Or algorithms that work in practice?