

CS394R

Reinforcement Learning: Theory and Practice

Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Morning Colleagues

- Are there any questions?

Logistics

- Make progress on final projects!

Logistics

- Make progress on final projects!
- This week's readings: financial applications

Logistics

- Make progress on final projects!
- This week's readings: financial applications
 - One old one, one newer

Logistics

- Make progress on final projects!
- This week's readings: financial applications
 - One old one, one newer
- Next week's readings: robot applications
 - **Kohl:** Policy gradient RL somewhat like Moody's RRL

Logistics

- Make progress on final projects!
- This week's readings: financial applications
 - One old one, one newer
- Next week's readings: robot applications
 - **Kohl:** Policy gradient RL somewhat like Moody's RRL
 - **Ng:** Helicopter learning: model-based, variance reduction

Logistics

- Make progress on final projects!
- This week's readings: financial applications
 - One old one, one newer
- Next week's readings: robot applications
 - **Kohl:** Policy gradient RL somewhat like Moody's RRL
 - **Ng:** Helicopter learning: model-based, variance reduction
 - **Riedmiller:** Learning directly from vision

Class Discussion

- Hassaan

Kearns on RL in Finance

- This paper wasn't tested against state of the art methods such as VWAP strategies. That could have led to something fieldable

Kearns on RL in Finance

- This paper wasn't tested against state of the art methods such as VWAP strategies. That could have led to something fieldable
- Dark pool algorithm implemented, may still be used

Kearns on RL in Finance

- This paper wasn't tested against state of the art methods such as VWAP strategies. That could have led to something fieldable
- Dark pool algorithm implemented, may still be used
- Biggest promise: layering on top of existing best practices