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

- First 3 programming assignments past due

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

  n
- Complete the class survey by Wed. at 5pm

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

  n
- Complete the class survey by Wed. at 5pm
- Literature review due November 10th

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

  n
- Complete the class survey by Wed. at 5pm
- Literature review due November 10th
  - At least 10 refs

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

  n
- Complete the class survey by Wed. at 5pm
- Literature review due November 10th
  - At least 10 refs
  - Also reiterate proposal; be more concrete; answer the biggest questions; preliminary results

Logistics

- First 3 programming assignments past due
  - Until literature survey is due: 1 point off
  - Until last day of class: 2 points off
  - After last day of class: 3 points off
  - After final project is due: 4 points off

  n
- Complete the class survey by Wed. at 5pm
- Literature review due November 10th
  - At least 10 refs
  - Also reiterate proposal; be more concrete; answer the biggest questions; preliminary results
  - Partial draft for final project

Overview

- RMax: model-based learning in polynomial time

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)
  - Q-learning vs. RMax (videos)

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)
  - Q-learning vs. RMax (videos)
- Rmax not built to be practical

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)
  - Q-learning vs. RMax (videos)
- Rmax not built to be practical
  - Built to be provably convergent

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)
  - Q-learning vs. RMax (videos)
- Rmax not built to be practical
  - Built to be provably convergent
- MBIE: more practical

Overview

- RMax: model-based learning in polynomial time
  - High-level idea (pdf)
  - Q-learning vs. RMax (videos)
- Rmax not built to be practical
  - Built to be provably convergent
- MBIE: more practical
- Fitted R-Max: Extend to continuous state space (pdf)

Discussion

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

Discussion

- What's more interesting? Sample complexity? Or average loss?