

CS394R

Reinforcement Learning: Theory and Practice

Peter Stone

Department of Computer Science
The University of Texas at Austin

Good Morning Colleagues

Good Morning Colleagues

- Are there any questions?

Logistics

Logistics

- Registering for the course (roster)

Logistics

- Registering for the course (roster)
- Nice responses!

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza
 - default: 9/10

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza
 - default: 9/10
 - Look for programming assignment opportunities

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza
 - default: 9/10
 - Look for programming assignment opportunities
 - First due Thursday

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza
 - default: 9/10
 - Look for programming assignment opportunities
 - First due Thursday
 - I have author's responses to exercises (1st edition)

Logistics

- Registering for the course (roster)
- Nice responses!
 - Length and content mostly good
 - Be clear and specific
 - Short and focussed is fine
 - Help me help you
 - Won't answer everything: ask in class or on piazza
 - default: 9/10
 - Look for programming assignment opportunities
 - First due Thursday
 - I have author's responses to exercises (1st edition)
- If you missed last Thursday . . .

My Role

My Role

- My role isn't to teach RL

My Role

- My role isn't to teach RL
- It's to help **you** learn RL

My Role

- My role isn't to teach RL
- It's to help **you** learn RL
 - provide context
 - guide your learning (assign readings, activities)
 - clarify misconceptions

My Role

- My role isn't to teach RL
- It's to help **you** learn RL
 - provide context
 - guide your learning (assign readings, activities)
 - clarify misconceptions
- You have to do the learning

My Role

- My role isn't to teach RL
- It's to help **you** learn RL
 - provide context
 - guide your learning (assign readings, activities)
 - clarify misconceptions
- You have to do the learning
- Read, write, present, ask, answer, program (investigate)

Logistics

Logistics

- Programming assignment due Thursday
- Next readings posted

Logistics

- Programming assignment due Thursday
- Next readings posted
 - **2nd edition!!**

Logistics

- Programming assignment due Thursday
- Next readings posted
 - **2nd edition!!**
 - MDPs and Dynamic programming

Logistics

- Programming assignment due Thursday
- Next readings posted
 - **2nd edition!!**
 - MDPs and Dynamic programming
 - Mostly chapter 3 Tuesday, then chapter 4

Logistics

- Programming assignment due Thursday
- Next readings posted
 - **2nd edition!!**
 - MDPs and Dynamic programming
 - Mostly chapter 3 Tuesday, then chapter 4
 - Discussion leaders

Logistics

- Programming assignment due Thursday
- Next readings posted
 - **2nd edition!!**
 - MDPs and Dynamic programming
 - Mostly chapter 3 Tuesday, then chapter 4
 - Discussion leaders
- Look at resources page

Let's Play!

Let's Play!

- I'm a 2-armed bandit

Let's Play!

- I'm a 2-armed bandit
- As a class, you choose which arm: 3 times around.

Let's Play!

- I'm a 2-armed bandit
- As a class, you choose which arm: 3 times around.
- Maximize your payoff.

Let's Play!

- I'm a 2-armed bandit
- As a class, you choose which arm: 3 times around.
- Maximize your payoff.
- The answer:

Let's Play!

- I'm a 2-armed bandit
- As a class, you choose which arm: 3 times around.
- Maximize your payoff.
- The answer:

```
(defun l () (+ 5 (random 7)))
```

```
(defun r ()  
  (let ((x (random 3)))  
    (case x  
      (0 20)  
      (1 0)  
      (2 (+ 7 (random 11))))  
    )))
```

- What about minimizing risk?

N-armed bandit in practice?

N-armed bandit in practice?

- Choosing mechanics
- Choosing a barber/hairdresser

Student-led Discussion

- Aishwarya Padmakumar on non-stationarity

Assignments

- Join piazza!

Assignments

- Join piazza!
- 1st programming assignment due Thursday
- Read Chapters 3 and 4

Assignments

- Join piazza!
- 1st programming assignment due Thursday
- Read Chapters 3 and 4
- Send a reading response by 5pm Monday

Assignments

- Join piazza!
- 1st programming assignment due Thursday
- Read Chapters 3 and 4
- Send a reading response by 5pm Monday
- Need a discussion leader volunteer