

CS325 Artificial Intelligence

Markov Decision Process Review

Cengiz Günay, Emory Univ.

(AN UNMATCHED LEFT PARENTHESIS
CREATES AN UNRESOLVED TENSION
THAT WILL STAY WITH YOU ALL DAY.

Spring 2013

Midterm Over, Let's Review Some and Have Fun

CajunBot I and II:



Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**



Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 1$, $\gamma = 1$, cost = -4

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a	84	88	92	96
b	-100	92	96	+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 1$, $\gamma = 1$, cost = -4

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Value of a_4 after one iteration?

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				76
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Value of a_4 after one iteration?

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Value of a_4 after it converges?

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				95
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Value of a_4 after it converges?

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a				
b	-100			+100

Movement probability:

- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Optimal policy after convergence?

Markov Decision Processes (MDPs) Exercises

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s).$$

	1	2	3	4
a	→	→	↓	↓
b	-100	↓	→	+100

Movement probability:


- p succeeds
- $(1 - p)$ fails, go reverse

Question: **Find final values of all positions.**

- For $p = 0.8$, $\gamma = 1$, cost = -4
- Optimal policy after convergence?

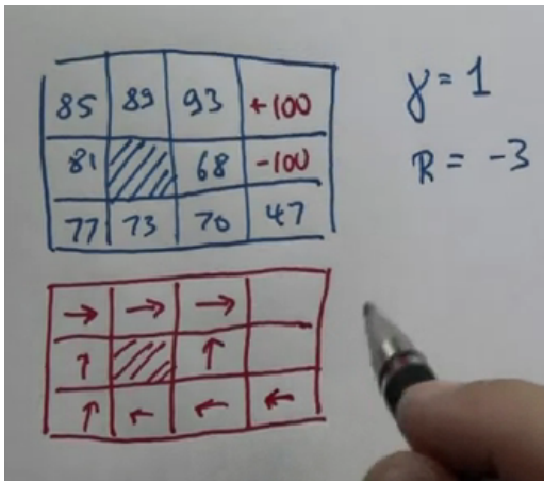
What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s), \quad P(s, a) = 0.8$$




What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a)V(s') \right] + R(s), \quad P(s, a) = 0.8$$



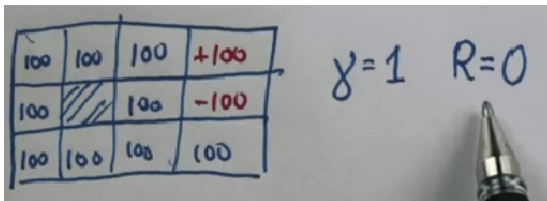
What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s), \quad P(s, a) = 0.8$$



What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s), \quad P(s, a) = 0.8$$




100	100	100	+100
100		100	-100
100	100	100	100

$\gamma = 1$ $R = 0$


What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a)V(s') \right] + R(s), \quad P(s, a) = 0.8$$




What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a) V(s') \right] + R(s), \quad P(s, a) = 0.8$$



What happens When Value Iterations Converge?

$$V(s) \leftarrow \left[\arg \max_a \gamma \sum_{s'} P(s'|a)V(s') \right] + R(s), \quad P(s, a) = 0.8$$



Exploration and Prediction with MDPs and Reinforcement Learning



Midterm Evaluations

Please post your evaluations! You have the chance to grade me! :)

Midterm Evaluations

Please post your evaluations! You have the chance to grade me! :)

- Any questions about midterm problems?

Midterm Evaluations

Please post your evaluations! You have the chance to grade me! :)

- Any questions about midterm problems?
- Have a Happy Spring Break!

