

“The rotten tree-trunk, until the very moment when the storm-blast breaks it in two, has all the appearance of might it ever had.”

Isaac Asimov, *Foundation*



CMPUT 365

Introduction to RL

Plan

- A note on the Markov property
- More exercises

Please, interrupt me at any time!



A note on the Markov property

Definition: We say that $(S_0, A_0, R_1, S_1, A_1, R_2, S_2, A_2, \dots)$ has the *Markov property* if for any $t \geq 0$, $s_0, s_1, \dots, s_{t+1} \in \mathcal{S}$, $a_0, a_1, \dots, a_t \in \mathcal{A}$, and $r_1, r_2, \dots, r_{t+1} \in \mathcal{R}$, it holds that

$$\Pr(R_{t+1} = r_{t+1}, S_{t+1} = s_{t+1} \mid S_0 = s_0, A_0 = a_0, R_1 = r_1, \dots, R_t = r_t, S_t = s_t, A_t = a_t)$$

only depends on the values of s_t , a_t , s_{t+1} and r_{t+1} . In particular, none of the other past values matter when calculating probabilities of the form above. That is:


$$\begin{aligned} &\Pr(R_{t+1} = r_{t+1}, S_{t+1} = s_{t+1} \mid S_0 = s_0, A_0 = a_0, R_1 = r_1, \dots, R_t = r_t, S_t = s_t, A_t = a_t) \\ &= \Pr(R_{t+1} = r_{t+1}, S_{t+1} = s_{t+1} \mid S_t = s_t, A_t = a_t). \end{aligned}$$

A note on the Markov property

Definition: We say that $(S_0, A_0, R_1, S_1, A_1, R_2, S_2, A_2, \dots)$ has the *Markov property* if for any $t \geq 0$, $s_0, s_1, \dots, s_{t+1} \in \mathcal{S}$, $a_0, a_1, \dots, a_t \in \mathcal{A}$, and $r_1, r_2, \dots, r_{t+1} \in \mathcal{R}$, it holds that

The Markov property does not mean that the state representation tells all that would be useful to know, only that it has not forgotten anything that would be useful to know.


$$\begin{aligned} \Pr(R_{t+1} = r_{t+1}, S_{t+1} = s_{t+1} \mid S_0 = s_0, A_0 = a_0, R_1 = r_1, \dots, R_t = r_t, S_t = s_t, A_t = a_t) \\ = \Pr(R_{t+1} = r_{t+1}, S_{t+1} = s_{t+1} \mid S_t = s_t, A_t = a_t). \end{aligned}$$



Example

Example 3.4: Pole-Balancing

The objective in this task is to apply forces to a cart moving along a track so as to keep a pole hinged to the cart from falling over: A failure is said to occur if the pole falls past a given angle from vertical or if the cart runs off the track. The pole is reset to vertical after each failure. This task could be treated as episodic, where the natural



episodes are the repeated attempts to balance the pole. The reward in this case could be +1 for every time step on which failure did not occur, so that the return at each time would be the number of steps until failure. In this case, successful balancing forever would mean a return of infinity. Alternatively, we could treat pole-balancing as a continuing task, using discounting. In this case the reward would be -1 on each failure and zero at all other times. The return at each time would then be related to $-\gamma^{K-1}$, where K is the number of time steps before failure (as well as to the times of later failures). In either case, the return is maximized by keeping the pole balanced for as long as possible. ■

Exercise 3.6 of the Textbook

Exercise 3.6 Suppose you treated pole-balancing as an episodic task but also used discounting, with all rewards zero except for -1 upon failure. What then would the return be at each time? How does this return differ from that in the discounted, continuing formulation of this task? □

Solution Exercise 3.6 of the Textbook

Exercise 3.7 of the Textbook

Exercise 3.7 Imagine that you are designing a robot to run a maze. You decide to give it a reward of +1 for escaping from the maze and a reward of zero at all other times. The task seems to break down naturally into episodes—the successive runs through the maze—so you decide to treat it as an episodic task, where the goal is to maximize expected total reward (3.7). After running the learning agent for a while, you find that it is showing no improvement in escaping from the maze. What is going wrong? Have you effectively communicated to the agent what you want it to achieve?


Solution Exercise 3.7 of the Textbook

Exercise 10 from Quiz on Coursera

10. Imagine, an agent is in a maze-like gridworld. You would like the agent to find the goal, as quickly as possible. You give the agent a reward of +1 when it reaches the goal and the discount rate is 1.0, because this is an episodic task. When you run the agent its finds the goal, but does not seem to care how long it takes to complete each episode. How could you fix this? (**Select all that apply**)

- Give the agent a reward of 0 at every time step so it wants to leave.
- Set a discount rate less than 1 and greater than 0, like 0.9.
- Give the agent a reward of +1 at every time step.
- Give the agent -1 at each time step.

Practice Exercise



	1	2	3
4	5	6	7
8	9	10	11
12	13	14	

$R_t = -1$
on all transitions

$$p(6, -1 | 5, \text{right}) =$$

$$p(7, -1 | 7, \text{right}) =$$

$$p(10, r | 5, \text{right}) =$$

Practice Exercise – Modeling

Assume you have a bandit problem with 4 actions, where the agent can see rewards from the set $\mathcal{R} = \{-3.0, -0.1, 0, 4.2\}$. Assume you have the probabilities for rewards for each action: $p(r|a)$ for $a \in \{1, 2, 3, 4\}$ and $r \in \{-3.0, -0.1, 0, 4.2\}$. How can you write this problem as an MDP? Remember that an MDP consists of $(\mathcal{S}, \mathcal{A}, \mathcal{R}, P, \gamma)$.

More abstractly, recall that a Bandit problem consists of a given action space $\mathcal{A} = \{1, \dots, k\}$ (the k arms) and the distribution over rewards $p(r|a)$ for each action $a \in \mathcal{A}$. Specify an MDP that corresponds to this Bandit problem.

Practice Exercise – Modeling

Exercise 3.10 of the Textbook

Exercise 3.10 Prove the second equality in (3.10).

$$G_t = \sum_{k=0}^{\infty} \gamma^k = \frac{1}{1 - \gamma}$$

Solution Exercise 3.10 of the Textbook

Practice Exercise

Prove that the discounted sum of rewards is always finite, if the rewards are bounded: $|R_{t+1}| \leq R_{\max}$ for all t for some finite $R_{\max} > 0$.

$$|\sum_{i=0}^{\infty} \gamma^i R_{t+1+i}| < \infty \text{ for } \gamma \in [0, 1). \quad \text{Hint: Recall that } |a + b| < |a| + |b|.$$

Solution Practice Exercise

Next class

- What **I** plan to do:
 - Non-comprehensive overview of Value Functions & Bellman Equations (2nd half of Chapter 3).
- What I recommend **YOU** to do for next class:
 - Read Chapter 3, §3.4–§3.8 (pp. 57–69).
 - Submit Practice Quiz for Fundamental of RL: Value functions & Bellman equations (Week 3).
 - Start Graded Quiz for Fundamental of RL: Value functions & Bellman equations (Week 3).