

Introduction to Artificial Intelligence (INFO8006)

Exercise session 7

Markov decision processes

A **Markov decision process** is a tuple $(\mathcal{S}, \mathcal{A}, P, R)$ such that:

- \mathcal{S} is a set of **states** s ;
- \mathcal{A} is a set of **actions** a ;
- P is a (stationary) **transition model** such that $P(s' | s, a)$ denotes the probability of reaching state s' if action a is done in state s ;
- R is **reward function** that maps immediate (finite) reward values $R(s)$ obtained in states s .

Bellman equation

The **utility** of a state is the immediate reward for that state, plus the expected discounted utility of the next state, assuming that the agent chooses the optimal action:

$$V(s) = R(s) + \gamma \max_a \sum_{s'} P(s' | s, a) V(s').$$

Value/policy iteration

The **value iteration** algorithm provides a fixed-point iteration procedure for computing the state utilities $V(s)$. By denoting $V_i(s)$ the estimated utility at iteration i and starting from an initial guess $V_0(s)$, we update the estimates to be consistent with **Bellman equation**:

$$V_{i+1}(s) = R(s) + \gamma \max_a \sum_{s'} P(s' | s, a) V_i(s')$$

which always converges for $\gamma < 1$.

The **policy iteration** algorithm instead directly computes the **policy** π by alternating between **value evaluation** under the current **policy** estimate

$$V_i = V^{\pi_i}$$

and **policy** improvement

$$\pi_{i+1}(s) = \arg \max_a \sum_{s'} P(s' | s, a) V_i(s').$$

Reinforcement learning

Reinforcement learning is used in unknown MDP, i.e. MDP where the **transition model** and the **reward function** are unknown. The goal is still to extract the optimal **policy** by observing or interacting with the environment in order to jointly learn these dynamics and act upon them.

RL can be decomposed in **model-based** and **model-free** methods. The former estimates explicitly the environment and then use this estimation as the empirical MDP whereas the latter do not model the environment explicitly.

Temporal-difference learning

Temporal-difference learning consists in updating $V^\pi(s)$ each time the agent experiences a transition (s, r, a, s') under a policy π .

When a transition from s to s' occurs, the temporal-difference update steers $V^\pi(s)$ to better agree with the Bellman equations for a fixed policy, i.e.

$$V^\pi(s) \leftarrow V^\pi(s) + \alpha(r + \gamma V^\pi(s') - V^\pi(s))$$

where α is the **learning rate**.

Q-Learning

By defining **Q-values** as $V(s) = \max_a Q(s, a)$, we have

$$Q(s, a) = R(s) + \gamma \sum_{s'} P(s' | s, a) \max_{a'} Q(s', a').$$

The state-action-values $Q(s, a)$ can be learned in a model-free fashion using a temporal-difference method known as **Q-Learning**. Q-Learning consists in updating $Q(s, a)$ each time the agent experiences a transition. The update equation for TD Q-Learning is

$$Q(s, a) \leftarrow Q(s, a) + \alpha(r + \gamma \max_{a'} Q(s', a') - Q(s, a)).$$

In session exercises: Ex. 1, Ex. 4

Part I

Making decisions

Exercise 1 Micro-Blackjack (UC Berkeley CS188, Spring 2014)

In Micro-Blackjack, you repeatedly draw a card (with replacement) that is likely to be a 2, 3 or 4. At each step, if the total score of the cards is lower than 6, you can either “draw” (d) or “cash” (c). Otherwise, you can only cash. When you cash, the game stops and your utility is equal to your total score (up to 5) plus 1, or zero if you get a total of 6 or higher. Until you cash and after it, you receive no reward. There is no discount ($\gamma = 1$).

1. Formalize Micro-Blackjack as an MDP.
2. Derive the optimal policy for this MDP.

Exercise 2 Crying Baby Problem

To earn some money you have decided to do babysitting. Tonight you will take care of a 9 months old baby. His parents asked you to feed him when he is hungry. Because you are a super-babysitter you know from your experience that the probability of the baby crying when he is hungry is equal to 0.8 and to cry when he is not hungry is equal to 0.1. The cost of feeding the baby is equal to 5 \$ whereas if you don't feed him while he is hungry it costs you 10 \$ (because the parent will be upset and will reduce your pay). You will have many things to do tonight and so you decide to optimize your time by just checking every hour whether the baby is crying. Between two intervals of time you assume that the baby has a probability of 0.2 to get hungry. You can also assume that the baby has a uniform probability of being hungry when you start your babysitting.

1. Draw the dynamic Bayesian network as well as the conditional probability tables associated with your babysitting of tonight.
2. Give the sequence of belief states if you assume the sequence of observation-action pairs $(0, 0)$, $(1, 1)$, $(0, 0)$, $(0, 0)$.
3. From the sequence of belief states, determine the sequence of expected rewards.

Exercise 3 Pursuit Evasion (UC Berkeley CS188, Spring 2014)

Pacman is trapped in the following 2 by 2 maze with a hungry ghost. When it is his turn to move, Pacman must move one step horizontally or vertically to a neighboring square. When it is the ghost's turn, he must also move one step horizontally or vertically. The ghost and Pacman alternate moves. After every move (by either the ghost or Pacman), if Pacman and the ghost occupy the same square, Pacman is eaten and receive utility -100 . Otherwise, he receives utility of 1. The ghost attempts to minimize the utility that Pacman receives. Pacman makes the first move. The game is not guaranteed to terminate.



For example, with a discount factor of $\gamma = 1$, if Pacman moves right, ghost moves down, then Pacman moves left, Pacman earns a reward of 1 initially ($R(s_0)$) and after the two first moves and -100 after the last move, for a total utility of -97 .

1. Assume a discount factor $\gamma = 0.5$, where the discount factor is applied once every time either Pacman or the ghost moves. What is the Minimax value of the complete (infinite) game?
2. Why are the value/policy iteration algorithms superior to Minimax for solving this game?

This game is similar to an MDP because rewards are earned at every timestep. However, it is also an adversarial game involving decisions by two agents. Let s be the state (*e.g.* the positions of Pacman and the ghost and who is playing) and let $A(s)$ be the set of actions available to the player in state s . Let $s' = T(s, a)$ denote the successor of a state s resulting from the action a . Finally, let $R(s)$ denote the utility received after moving to state s .

4. Write down an expression (analogous to Bellman equation) for $V(s)$, the value of the current state if it is Pacman's turn.

Part II

Reinforcement learning

Exercise 4 Q-learning


An agent is in an unknown environment where there are three states $\{A, B, C\}$ and two actions $\{0, 1\}$. We are given the following tuples (s, a, r, s') , generated by taking actions in the environment.

s	a	r	s'
A	0	+2	A
C	1	-2	A
B	1	+1	B
A	0	-1	B
B	1	-2	C
C	0	+4	B
B	0	+1	A

Assuming a discount factor $\gamma = 0.5$ and a learning rate $\alpha = 0.75$,

1. Apply the Q -learning algorithm to obtain state-action-value $Q(s, a)$ estimates. Estimates are initialized to 0.
2. We now switch to a feature-based estimator $\hat{Q}(s, a) = w_0 + w_1 f_1(s, a)$, with $f_1(s, a) = 2a - 1$. Starting from weights $w_0 = w_1 = 0$, update the weights according to the approximate Q -learning algorithm.

Exercise 5 Passive RL

3	-6	-1	+5
2			
1		-8	+3
	1	2	3

Consider the grid-world given above and an agent who is trying to learn the optimal policy. The agent starts from the bottom-left corner and can take the actions **north** (N), **south** (S), **west** (W) and **east** (E). Rewards are only awarded for reaching the terminal (shaded) states. You observe the following trials, whose trajectories are sequences of tuples $(s_t^i, r_t^i, a_t^i, s_{t+1}^i)$.

t	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5
0	(1, 1), 0, N , (1, 2)	(1, 1), 0, N , (1, 2)	(1, 1), 0, N , (1, 2)	(1, 1), 0, N , (1, 2)	(1, 1), 0, N , (1, 2)
1	(1, 2), 0, E , (2, 2)	(1, 2), 0, E , (2, 2)	(1, 2), 0, E , (2, 2)	(1, 2), 0, E , (2, 2)	(1, 2), 0, E , (2, 2)
2	(2, 2), 0, N , (2, 3)	(2, 2), 0, E , (3, 2)	(2, 2), 0, S , (2, 1)	(2, 2), 0, E , (3, 2)	(2, 2), 0, E , (3, 2)
3	(2, 3), -1, \emptyset , \emptyset	(3, 2), 0, N , (3, 3)	(2, 1), -8, \emptyset , \emptyset	(3, 2), 0, W , (2, 2)	(3, 2), 0, S , (3, 1)
4		(3, 3), +5, \emptyset , \emptyset		(2, 2), 0, N , (2, 3)	(3, 1), +3, \emptyset , \emptyset
5				(2, 3), -1, \emptyset , \emptyset	

Assuming a discount factor $\gamma = 1$,

1. Perform direct utility estimation of the expected utilities $V^\pi(s)$, given the four first trials.
2. Update the estimated expected utilities with respect to the fifth trial using temporal-difference learning. Assume a learning rate $\alpha = 0.5$.

Exercise 6 Football

ULiège's football team is playing against UCL's team next week. With a lot of losses this season, Liège needs to improve their attack strategy to win the game and increase their popularity. Luckily, the team captain follows INFO8006 and knows how to model the attack as a Markov Decision Process. The captain considers four states **close** (C), **away** (A), **fail** (F), and **goal** (G), and two actions **pass** (P) and **shoot** (S). Although the transition probabilities are unsure, the possible transitions (s, a, s') are known. To each transition is associated an increase/decrease of the team's popularity.

s	a	s'	$R(s, a, s')$
C	P	C	+1
C	P	A	-1
C	P	F	-2
C	S	C	+3
C	S	F	-5
C	S	G	+10
A	P	C	+2
A	P	A	0
A	P	F	-3
A	S	C	+3
A	S	F	-10
A	S	G	+20

The current strategy of the team is to always shoot. Last match, they had several attack opportunities, resulting in the following actions.

s	C	C	C	C	A	A	A	A
a	S	S	S	S	S	S	S	S
s'	G	C	G	F	F	C	F	F

Assuming a discount factor $\gamma = 0.75$ and a learning rate $\alpha = 0.25$,

1. Build an estimator of the transition model $P(s'|s, a)$ and, from it, determine the expected utility V^π of the team's current policy π .

The captain found the tapes of the previous season where they had much more success. Together with the team, the captain selects the following instructive actions.

s	C	C	A	A	C	A	A	C
a	S	S	S	P	P	P	S	P
s'	G	C	F	C	C	C	F	F

2. Given the selected tuples, apply the Q -learning algorithm to obtain state-action-value $Q(s, a)$ estimates. Estimates are initialized to 0.
3. Determine the optimal policy according to the state-action-value estimates.

Supplementary exercises

Exercise 7 (AIMA, Ex 16.15)

Consider a student, Sam, who has the choice to buy or not buy a textbook for a course. We'll model this as a decision problem with one Boolean decision node, B , indicating whether the agent chooses to buy the book, and two Boolean chance nodes, M , indicating whether the student has mastered the material of the course, and F , indicating whether the student fails the course. There is also a utility node, U . Sam has an additive utility function: 0 for not buying the book and -100 \$ for buying it; and 0 for failing the course and 2000 \$ for passing it. Sam's conditional probability estimates are as follows:

B	M	$P(M B)$	$P(F = 1 B, M)$
0	0	0.3	0.7
1	0	0.1	0.5
0	1	0.7	0.2
1	1	0.9	0.1

1. Draw the decision network for this problem.
2. Compute the expected utility of buying the book and of not buying it.
3. What should Sam do?

Exercise 8 (AIMA, Ex 17.6)

The slides of the course states that the Bellman operator

$$(BV)(s) = R(s) + \gamma \max_{a \in A(s)} \sum_{s'} P(s'|s, a) V(s')$$

is a contraction.

1. Show that, for any two functions $f : \mathbb{R} \mapsto \mathbb{R}$ and $g : \mathbb{R} \mapsto \mathbb{R}$,

$$|\max_x f(x) - \max_x g(x)| \leq \max_x |f(x) - g(x)|.$$

2. Apply the previous result to prove that the Bellman operator is a contraction of factor γ , that is

$$\|(BV) - (BV')\|_\infty \leq \gamma \|V - V'\|_\infty.$$

The infinity norm of a function $f : \mathbb{R} \mapsto \mathbb{R}$ is defined as

$$\|f\|_\infty = \sup_x |f(x)|.$$

Exercise 9 (AIMA, Ex 17.8)

Consider the 3×3 world shown in Figure 1. The transition model is the same as in the 4×3 2: 80% of the time the agent goes in the direction it selects; the rest of the time it moves at right angles to the intended direction.




Figure 1. (a) 3×3 world for Exercise 9. The reward for each state is indicated. The upper right square is a terminal state. (b) 101×3 world for Exercise 10. The start state has reward 0.




Figure 2. (a) A simple 4×3 environment that presents the agent with a sequential decision problem. The two terminal states have reward $+1$ and -1 , respectively, and all other states have a reward of -0.04 . (b) Illustration of the transition model of the environment: the “intended” outcome occurs with probability 0.8, but with probability 0.2 the agent moves at right angles to the intended direction. A collision with a wall results in no movement.

Implement value iteration for this world for each value of r below. Use discounted rewards with a discount factor $\gamma = 0.99$. Show the policy obtained in each case. Explain intuitively why the value of r leads to each policy.

1. $r = -100$
2. $r = -3$
3. $r = 0$
4. $r = +3$

Exercise 10 (AIMA, Ex 17.9)

Consider the 101×3 world shown in Figure 1. In the start state the agent has a choice of two deterministic actions, **up** or **down**, but in the other states the agent has one deterministic action, **right**. Assuming a discounted reward function, compute the utility of each action as a function of γ . For what values of the discount factor γ should the agent choose **up** and for which **down**?

This simple example actually reflects many real-world situations in which one must weigh the value of an immediate action versus the potential long-term consequences, such as choosing to dump pollutants into a lake.

Exercise 11 (AIMA, Ex 17.10)

Consider an undiscounted MDP having three states, $(1, 2, 3)$, with rewards $-1, -2, 0$, respectively. State 3 is a terminal state. In states 1 and 2 there are two possible actions: a and b . The transition model is as follows:

- In state 1, action a moves the agent to state 2 with probability 0.8 and makes the agent stay put with probability 0.2.
- In state 2, action a moves the agent to state 1 with probability 0.8 and makes the agent stay put with probability 0.2.
- In either state 1 or state 2, action b moves the agent to state 3 with probability 0.1 and makes the agent stay put with probability 0.9.

Answer the following questions:

1. What can be determined *qualitatively* about the optimal policy in states 1 and 2?
2. Apply policy iteration, showing each step in full, to determine the optimal policy and the values of states 1 and 2. Assume that the initial policy has action b in both states.
3. What happens to policy iteration if the initial policy has action a in both states? Does discounting help? Does the optimal policy depend on the discount factor?

Exercise 12 (AIMA, Ex 17.19)

A Dutch auction is similar to an English auction, but rather than starting the bidding at a low price and increasing, the seller starts at a high price and gradually lowers the price until some buyer is willing to accept that price. If multiple bidders accept the price, one is arbitrarily chosen as the winner. More formally, the seller begins with a price p and gradually lowers p by increments of d until at least one buyer accepts the price.

Assuming all bidders act rationally, is it true that for arbitrarily small d , a Dutch auction will always result in the bidder with the highest value for the item obtaining the item? If so, show mathematically why. If not, explain how it may be possible for the bidder with highest value for the item not to obtain it.

Exercise 13 (AIMA, Ex 17.21)

Teams in the National Hockey League historically received 2 points for winning a game and 0 for losing. If the game is tied, an overtime period is played; if nobody wins in overtime, the game is a tie and each team gets 1 point. But league officials felt that teams were playing too conservatively in overtime (to avoid a loss), and it would be more exciting if overtime produced a winner. So in 1999 the officials experimented in mechanism design: the rules were changed, giving a team that loses in overtime 1 point, not 0. It is still 2 points for a win and 1 for a tie.

1. Was hockey a zero-sum game before the rule change? After?
2. Suppose that at a certain time t in a game, the home team has probability p of winning in regulation time, probability $0.78 - p$ of losing, and probability 0.22 of going into overtime, where they have probability q of winning, $0.9 - q$ of losing, and 0.1 of tying. Give equations for the expected value for the home and visiting teams.
3. Imagine that it were legal and ethical for the two teams to enter into a pact where they agree that they will skate to a tie in regulation time, and then both try in earnest to win in overtime. Under what conditions, in terms of p and q , would it be rational for both teams to agree to this pact?
4. Some experts report that since the rule change, the percentage of games with a winner in overtime went up 18.2%, as desired, but the percentage of overtime games also went up 3.6%. What does that suggest about possible collusion or conservative play after the rule change?

Quiz

Which of the following is true? In Markov Decision Processes, ...

- the closer the discount factor γ to 0, the higher the utility of future rewards.
- the closer the discount factor γ to 0, the longer Value Iteration may take to converge.
- the closer the discount factor γ to 1, the greedier the optimal agent.
- the closer the discount factor γ to 1, the longer Value Iteration may take to converge.

Q-Learning ...

- is a model-based reinforcement learning algorithm.
- is an on-policy reinforcement learning algorithm.
- converges to an optimal policy, but only when acting optimally.
- requires a random exploration strategy to converge to an optimal policy.

Assume that we run ϵ -greedy Q-learning until convergence. What is the optimal policy π^* we obtain for an arbitrary state s ?

- $\pi^*(s) = \arg \max_s V(s)$
- $\pi^*(s) = \begin{cases} \arg \max_a Q(s, a) & \text{with probability } 1 - \epsilon \\ \text{random action} & \text{with probability } \epsilon \end{cases}$
- $\pi^*(s) = \arg \max_a Q(s, a)$
- $\pi^*(s) = \arg \max_s Q(s, a)$

Let us consider a robot wandering around at the Montefiore Institute. Which of the following is true?

- The robot and its environment can be modeled as a partially observable MDP.
- The Kalman filter can be used for determining accurately its past trajectory in the building.
- A convolutional neural network would be too large to fit in the robot's memory.
- Value iteration (with an admissible heuristic) can be used for decoding the speech of its visitors.

The Bellman equations $Q(s, a) = R(s) + \gamma \sum_{s'} P(s'|s, a) \max_{a'} Q(s', a')$ form a system of n non-linear equations with as many unknowns. Which of the following is true?

- n is the size of the action space \mathcal{A} .
- n is the size of the state space \mathcal{S} .
- n is the size of the state-action space $\mathcal{S} \times \mathcal{A}$.
- n is the sum of the sizes of the action and state spaces.

In DQN (Mnih et al, 2015), a reinforcement learning agent is trained to ...

- to classify images.
- to control a robot in a simulated environment.
- to play Atari video games.
- to play the game of Go.

In DQN (Mnih et al, 2015), the Q-table is approximated with ...

- a hash table.
- a transposition table.
- a linear regression model.
- a convolutional neural network.