

Solving MDPs

Here, Transition Probabilities and Reward Functions are known, and we want to find the optimal policy for the agent to maximize its discounted total utility. There are two ways:

1. Value Iteration with Policy Extraction

Repeat Value Iteration until convergence, then run Policy Extraction.

Value Iteration: $\forall s \in S, V_{k+1}(s) \leftarrow \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V_k(s')]$

Policy Extraction: $\forall s \in S, \pi^*(s) = \operatorname{argmax}_a Q^*(s, a) = \operatorname{argmax}_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V^*(s')]$

2. Policy Iteration

Repeatedly perform Policy Evaluation then Policy Improvement.

Policy Evaluation:

Primary Method: $V^\pi(s) = \sum_{s'} T(s, \pi(s), s') [R(s, \pi(s), s') + \gamma V^\pi(s')]$

Solve the system of equations for each $V^\pi(s)$.

Alternate Method: $V_{k+1}^{\pi_i}(s) \leftarrow \sum_{s'} T(s, \pi_i(s), s') [R(s, \pi_i(s), s') + \gamma V_k^{\pi_i}(s')]$

Repeat until convergence.

Policy Improvement: $\pi_{i+1}(s) = \operatorname{argmax}_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma V^{\pi_i}(s')]$

Reinforcement Learning

In RL problems, T and R are not known. Instead, we observe episodes of the agent's behavior with several (s, a, s', r) samples.

Model-Based Learning

Create a model to approximate $\hat{T}(s, a, s')$ and $\hat{R}(s, a, s')$. We keep counts of the (s, a, s', r) samples observed, then calculate \hat{T} and \hat{R} based on these counts.

$$\hat{T}(s, a, s') = \frac{\text{count}(s, a, s')}{\text{count}(s, a)}, \text{ and } \hat{R}(s, a, s') = r$$

As the number of samples increases, we converge to the true T and R .

Model-Free Learning

On-policy methods try to learn the value of states according to a certain (fixed) policy. Off-policy methods try to learn other information that could be used to find a good policy.

Direct Evaluation (On-policy)

For each state, average out the discounted total reward from all the episodes that visited that state.

$$V^\pi(s) = \frac{\text{sum of discounted total rewards of episodes that visit } s}{\text{number of episodes that visit } s}$$

Temporal Difference Learning (On-policy)

Learn from every experience by incorporating sample information into an exponential moving average.

$$\begin{aligned} \text{sample}_k &= R(s, \pi(s), s') + \gamma V_{k-1}^\pi(s') \\ V_k^\pi(s) &\leftarrow (1 - \alpha)V_{k-1}^\pi(s) + \alpha \cdot \text{sample}_k \end{aligned}$$

Q-Learning (Off-policy)

Q-value Iteration. $Q_{k+1}(s, a) \leftarrow \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma \max_{a'} Q_k(s', a')]$

Incorporate samples into exponential moving average:

$$\begin{aligned} \text{sample} &= R(s, a, s') + \gamma \max_{a'} Q(s', a') \\ Q(s, a) &\leftarrow (1 - \alpha)Q(s, a) + \alpha \cdot \text{sample} \end{aligned}$$

Approximate Q-learning (Off-policy)

Approximate Q and V functions as linear combinations of features.

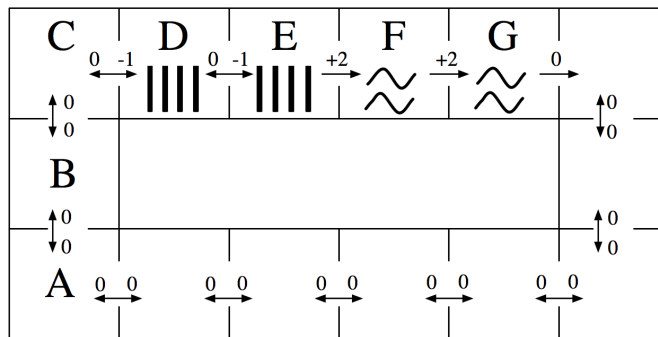
$$\begin{aligned} V(s) &= w_1 \cdot f_1(s) + w_2 \cdot f_2(s) + \dots + w_n \cdot f_n(s) = \vec{w} \cdot \vec{f}(s) \\ Q(s, a) &= w_1 \cdot f_1(s, a) + w_2 \cdot f_2(s, a) + \dots + w_n \cdot f_n(s, a) = \vec{w} \cdot \vec{f}(s, a) \end{aligned}$$

Update Rule:

$$\begin{aligned} \text{diff} &= [R(s, a, s') + \gamma \max_{a'} Q(s', a')] - Q(s, a) \\ w_i &\leftarrow w_i + \alpha \cdot \text{diff} \cdot f_i(s, a) \end{aligned}$$

1 MDPs: Grid-World Water Park

Consider the MDP drawn below. The state space consists of all squares in a grid-world water park. There is a single waterslide that is composed of two ladder squares and two slide squares (marked with vertical bars and squiggly lines respectively). An agent in this water park can move from any square to any neighboring square, unless the current square is a slide in which case it must move forward one square along the slide. The actions are denoted by arrows between squares on the map and all deterministically move the agent in the given direction. The agent cannot stand still: it must move on each time step. Rewards are also shown below: the agent feels great pleasure as it slides down the water slide (+2), a certain amount of discomfort as it climbs the rungs of the ladder (-1), and receives rewards of 0 otherwise. The time horizon is infinite; this MDP goes on forever.



- How many (deterministic) policies π are possible for this MDP?
- Fill in the blank cells of this table with values that are correct for the corresponding function, discount, and state. *Hint: You should not need to do substantial calculation here.*

	γ	$s = A$	$s = E$
$V_3(s)$	1.0		
$V_{10}(s)$	1.0		
$V_{10}(s)$	0.1		
$Q_1(s, \text{west})$	1.0	---	
$Q_{10}(s, \text{west})$	1.0	---	
$V^*(s)$	1.0		
$V^*(s)$	0.1		

- (c) Fill in the blank cells of this table with the Q-values that result from applying the Q-update for the transition specified on each row. You may leave Q-values that are unaffected by the current update blank. Use discount $\gamma = 1.0$ and learning rate $\alpha = 0.5$. Assume all Q-values are initialized to 0. (Note: the specified transitions would not arise from a single episode.)

	$Q(D, \text{west})$	$Q(D, \text{east})$	$Q(E, \text{west})$	$Q(E, \text{east})$
Initial:	0	0	0	0
Transition 1: ($s = D, a = \text{east}, r = -1, s' = E$)				
Transition 2: ($s = E, a = \text{east}, r = +2, s' = F$)				
Transition 3: ($s = E, a = \text{west}, r = 0, s' = D$)				
Transition 4: ($s = D, a = \text{east}, r = -1, s' = E$)				

2 Pacman with Feature-Based Q-Learning

We would like to use a Q-learning agent for Pacman, but the size of the state space for a large grid is too massive to hold in memory. To solve this, we will switch to feature-based representation of Pacman's state.

- (a) We will have two features, F_g and F_p , defined as follows:

$$F_g(s, a) = A(s) + B(s, a) + C(s, a)$$
$$F_p(s, a) = D(s) + 2E(s, a)$$

where

- $A(s)$ = number of ghosts within 1 step of state s
- $B(s, a)$ = number of ghosts Pacman touches after taking action a from state s
- $C(s, a)$ = number of ghosts within 1 step of the state Pacman ends up in after taking action a
- $D(s)$ = number of food pellets within 1 step of state s
- $E(s, a)$ = number of food pellets eaten after taking action a from state s

For this pacman board, the ghosts will always be stationary, and the action space is $\{left, right, up, down, stay\}$.



Calculate the features for the actions $\in \{left, right, up, stay\}$ from the current state shown in the figure.

- (b) After a few episodes of Q-learning, the weights are $w_g = -10$ and $w_p = 100$. Calculate the Q value for each action $\in \{left, right, up, stay\}$ from the current state shown in the figure.
- (c) We observe a transition that starts from the state above, s , takes action up , ends in state s' (the state with the food pellet above) and receives a reward $R(s, a, s') = 250$. The available actions from state s' are $down$ and $stay$. Assuming a discount of $\gamma = 0.5$, calculate the new estimate of the Q value for s based on this episode.
- (d) With this new estimate and a learning rate (α) of 0.5, update the weights for each feature.