

# Q-Learning

# Recap: Bellman Optimality

## Bellman Optimality

$$Q^*(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} Q^*(s', a'), \forall s, a$$

## VI: An iterative approach for estimating $Q^*$

$$Q \leftarrow \mathcal{T} Q$$

1. Need to know the transition
2. Only works for discrete small MDPs

# Recap: Bellman Optimality

**Q: if there is some  $Q(s, a)$ , such that the following holds:**

$$Q(s, a) = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} Q(s', a'), \forall s, a$$

**is this  $Q = Q^*$ ?**

# Today

Given MDP  $\mathcal{M} = (S, A, r, P, \gamma)$ ,

how to estimate  $Q^*(s, a), \forall s$  **WITHOUT** knowing  $P$   
i.e., how to learn  $Q^*$  (thus,  $\pi^*$ ) from experience

# Motivation

Computing a near-optimal policy to achieve the long-term goals w/o knowing or explicitly modeling the world



# Outline:

1. Q Learning

2. Revisit TD: Off-policy TD Learning

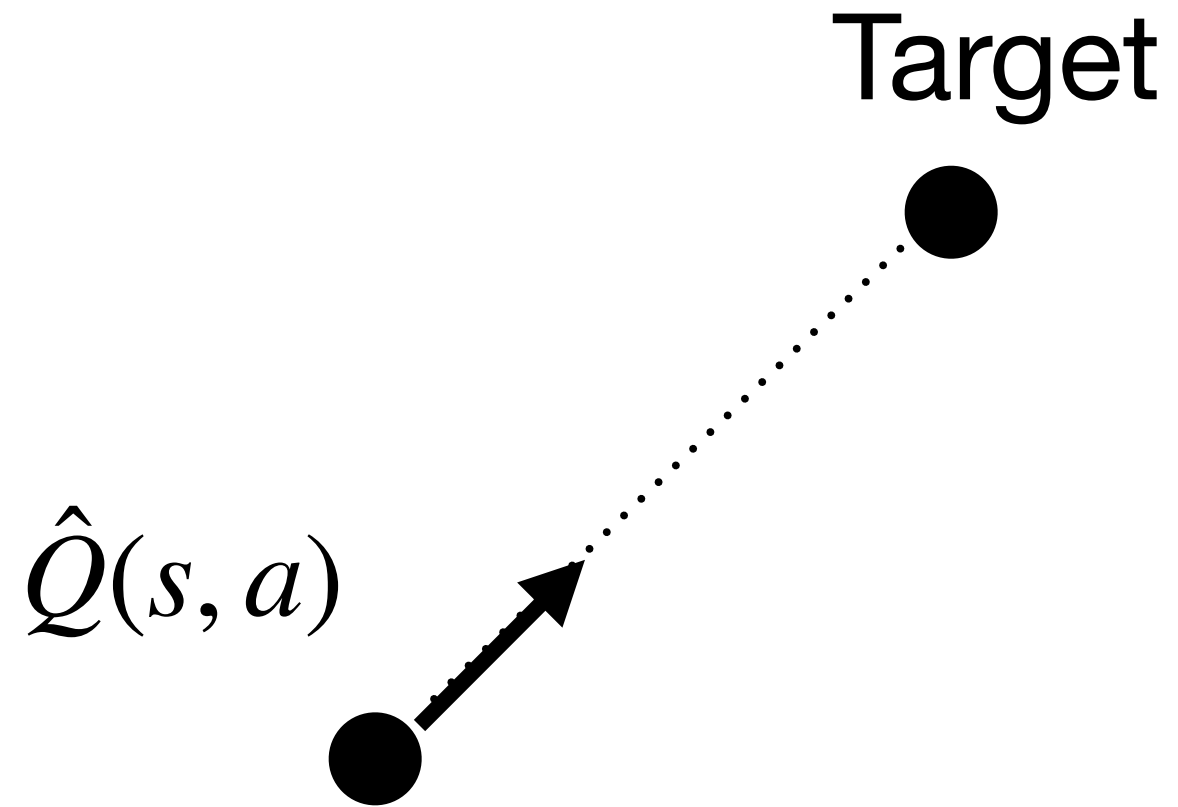

# Q Learning

$$Q^*(s, a) \approx r(s, a) + \gamma \max_{a'} \hat{Q}(s', a')$$

target

(Bootstrapping again)

$$\max_{a'} \hat{Q}(s', a')$$



Q-learning update: move to the target with a small step

# Q Learning

Given a **one-step transition**  $(s, a, r, s')$  where  $r = r(s, a)$ ,  $s' \sim P(\cdot | s, a)$ :

Q-learning updates the guess at  $(s, a)$  as follows:

$$\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \eta \left( \underbrace{r + \gamma \max_{a'} \hat{Q}(s', a')}_{\text{Q-Target}} - \hat{Q}(s, a) \right)$$

Q-Target  
(Constructed via Bootstrapping!)



# Q Learning

## How to collect data?

**Choice one**: trust current estimator  $\hat{Q}$ , always use  $\arg \max_a \hat{Q}(s, a)$

Issue: cannot explore (i.e., need to try something that hasn't been tried)

**Choice two (quite effective in practice)**:  $\epsilon$ -greedy

W/ prob  $\epsilon$ , select action uniform randomly

W/ prob  $1 - \epsilon$ , select greedy action  $\arg \max_a \hat{Q}(s, a)$

# TD Learning

Initialize  $\hat{Q}(s, a) = 0, \forall s, a$ . Set initial state  $s \in \mathcal{S}$

While True:

Take action  $a$  based on  $\epsilon$ -greedy of  $\hat{Q}$ , get reward  $r$  and next state  $s' \sim P(\cdot | s, a)$

Form Q-target  $r + \gamma \max_{a'} \hat{Q}(s', a')$

Update for  $s, a$ :  $\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \eta \left( r + \gamma \max_{a'} \hat{Q}(s', a') - \hat{Q}(s, a) \right)$

Set  $s \leftarrow s'$

# Interpret Q-learning as “SGD” on Bellman error

Q-learning is not the usual SGD, i.e., it is **not** running SGD to minimize a fixed loss function

Q-learning may be interpreted as running SGD on an **evolving** loss function (Bellman error)

$$\ell_{be}(\hat{Q}(s, a)) := \left( \hat{Q}(s, a) - y \right)^2, \text{ where } y = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} \hat{Q}(s', a')$$

$$\nabla \ell_{be}(x) |_{x=\hat{Q}(s, a)} := 2 \left( \hat{Q}(s, a) - y \right)$$

Unbiased  
estimate of  $y$

This keeps changing as  
we learning

$$\widetilde{\nabla} \ell_{be}(x) |_{x=\hat{Q}(s, a)} := 2 \left( \hat{Q}(s, a) - \left( r + \gamma \max_{a'} \hat{Q}(s', a') \right) \right)$$

# Q Learning Theory

[Informal] Assume the  $\epsilon$ -greedy strategy has non-trivial probability of visiting every state-action pair. Setting learning rate  $\eta$  properly, we will have:

$$\hat{Q}(s, a) \rightarrow Q^*(s, a), \forall s, a$$

when # of interactions approaches to  $\infty$

(concrete convergence rates are known as well)

# Demo: Q-learning on CartPole

Note: Cartpole's state is continuous, so we will need Q-learning w/ function approximation, e.g., neural network (we will get there very soon)

1. Does Q learning eventually learn a good policy
2. How does the  $\epsilon$  affect the learning

# Outline:

1. Q Learning

2. Revisit TD: Off-policy TD Learning

# TD Learning

Given  $(s, a, r, s')$ , where  $a \sim \pi(\cdot | s)$ ,  $s' \sim P(\cdot | s, a)$ , TD updates:

$$\hat{V}^\pi(s) \leftarrow \hat{V}^\pi(s) + \eta \left( r + \gamma \hat{V}^\pi(s') - \hat{V}^\pi(s) \right)$$

**On-policy:** data is generated from the policy  $\pi$  itself

**Off-policy:** data is generated from policy  $\pi_b$  where  $\pi_b \neq \pi$

Q: is Q-learning off-policy or on-policy?

# Motivation for off-policy evaluation

**Counterfactual:** what would happen if I did something different?



# Off-policy TD Learning

Setting: data is generated by  $\pi_b$ , but we want to estimate  $V^\pi$  for some  $\pi \neq \pi_b$

Key trick: importance weighting

$$\begin{aligned} V^\pi(s) &= \mathbb{E}_{a \sim \pi(\cdot|s)} \left( r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a)} V^\pi(s') \right) \\ &= \mathbb{E}_{a \sim \pi_b(\cdot|s)} \frac{\pi(a|s)}{\pi_b(a|s)} \left( r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot|s, a)} V^\pi(s') \right) \end{aligned}$$

Now action is  
sampled from  $\pi_b$

Importance weight

# Off-policy TD Learning

Given  $(s, a, r, s')$ , where  $a \sim \pi_b(\cdot | s)$ ,  $s' \sim P(\cdot | s, a)$ ,

Off-policy TD updates as follows:

$$\hat{V}^\pi(s) \leftarrow \hat{V}^\pi(s) + \eta \frac{\pi(a | s)}{\pi_b(a | s)} \left( r + \gamma \hat{V}^\pi(s') - V^\pi(s) \right)$$

Case 1:  $\pi(a | s)$  is large but  $\pi_b(a | s)$  is small

Case 2:  $\pi(a | s)$  is small but  $\pi_b(a | s)$  is large

# Off-policy TD Learning is SGD on TD loss

Given  $(s, a, r, s')$ , where  $a \sim \pi_b(\cdot | s)$ ,  $s' \sim P(\cdot | s, a)$ , Off-policy TD updates:

$$\hat{V}^\pi(s) \leftarrow \hat{V}^\pi(s) + \eta \frac{\pi(a | s)}{\pi_b(a | s)} \left( r + \gamma \hat{V}^\pi(s') - V^\pi(s) \right)$$

Check if it is doing one-step SGD on the TD loss:

$$\ell_{td}(\hat{V}^\pi(s)) = \left( \hat{V}^\pi(s) - y \right)^2 \text{ where } y = \mathbb{E}_{a \sim \pi(\cdot | s)} \left( r + \gamma \mathbb{E}_{s' \sim P(s, a)} \hat{V}^\pi(s') \right)$$

The off-policy TD update is one-step SGD on  $\ell_{td}$  (more in HW2)

# Summary

Q-Learning: online algorithm that learns  $Q^*$  (bootstrapping)

Exploration & Exploitation tradeoff:  $\epsilon$ -greedy is an effective heuristic

Off-policy policy evaluation: importance weighting  
(also known as inverse probability weighting in causal inference)