

230 2017-10-26

TD:  $v(S_t) \leftarrow v(S_t) + \alpha (R_t + \gamma v(S_{t+1}) - v(S_t))$  Tabular (one  $v$  entry per state  $s$ )  
 $w \leftarrow w + \alpha (R_t + \gamma w^T \phi(S_{t+1}) - w^T \phi(S_t))$  Linear Function Approximation  
 $w \leftarrow w + \alpha (R_t + \gamma v_w(S_{t+1}) - v_w(S_t)) \frac{\partial v_w(S_t)}{\partial w}$  General Function Approximation

TD error

Properties: Tabular: converges to  $V^\pi$  a.s. if  $\alpha$  decreased properly  
 Linear: converges to  $w_\infty$  s.t.  $MSE(w_\infty) \leq \frac{1}{1-\gamma} MSE(w^*)$   
 General: can diverge.

$\uparrow$   
 $\arg \min_w MSE(w)$

	DP	MC	TD
must know $P+R$	Y	N	N
must wait until episode end	N.A.	Y	N

An optimal estimator balances the bias/variance trade off.  $\leftarrow$

MC vs. TD

- What makes a better target,  $G_t$  or  $R_t + \gamma v(S_{t+1})$ ?
  - Each of these is an estimator of  $v^\pi(S_t)$
  - Mean Squared Error (MSE) is one measure of estimator quality
- For random  $X$  and  $\theta \in \mathbb{R}$ ,
- $$MSE(X) = E[(X - \theta)^2] = \underbrace{(E[X - \theta])^2}_{\text{bias}(X)^2} + \text{Var}(X)$$
- $$= \text{bias}(X)^2 + \text{Var}(X)$$

$$MSE(G_t) = \underbrace{\text{bias}(G_t)}_0^2 + \text{Var}(G_t)$$

$$\downarrow \text{Var}(R_t + \gamma R_{t+1} + \dots)$$

$$MSE(R_t + \gamma v(S_{t+1})) = \underbrace{\text{bias}(R_t + \gamma v(S_{t+1}))}_0^2 + \text{Var}(R_t + \gamma v(S_{t+1}))$$

$\downarrow$   
small when  $v$  is  
fairly accurate

MC vs TD

$$\hat{P}(s, a, s') = \frac{\#(s, a, s') \text{ transitions}}{\#(s, a) \text{ events}} \quad \hat{R}(s, a, s') = \text{mean}(r | s, a, s')$$

- These estimates of  $P+R$  maximize the likelihood of the observed data  
ML model of the MDP

- Given a fixed batch of data -  $(s, a, r, s')$  tuples


- If every state observed one or more times

- Then TD applied to convergence gives  $V_{\hat{P}, \hat{R}}^{\pi}$  if

$\hat{P} + \hat{R}$  were the transition + reward functions

Sarsa: Using TD for policy improvement/control

Idea: Use TD to estimate  $q^\pi$  & simultaneously change  $\pi$  to be greedy w.r.t to  $q^\pi$ .



View states as  $(s, a)$  pairs:  $(S_t, A_t)$   
 $((S_t, A_t), R_t, (S_{t+1}, A_{t+1})) \Rightarrow v(S_t, A_t)$   
estimates  $\mathbb{E}\left[\sum_{k=0}^{\infty} \gamma^k R_{t+k} \mid S_t=s, A_t=a, \pi\right]$

originally:  
 $(S_t, R_t, S_{t+1}) \Rightarrow v(S_t)$ , which  
estimates  $\mathbb{E}[\sum_{k=0}^{\infty} \gamma^k R_{t+k} \mid S_t=s, \pi]$

Updates:  
 $q(s, a) \leftarrow q(s, a) + \alpha (r + \gamma \underbrace{q(s', a')} - q(s, a))$  } TD for policy evaluation.  
Properties as before.

Can replace with  $\sum_{\bar{a}} \pi(s', \bar{a}) \cdot q(s', \bar{a})$  - but not  
necessary.

General form:

$$w \leftarrow w + \alpha (r + \gamma q_w(s', a') - q_w(s, a)) \frac{\partial q_w(s, a)}{\partial w}$$

## Control using TD (Sarsa):

Init:  $q(s, a) \leftarrow$  arbitrary

Repeat for each episode  $\epsilon$

$s \sim d_0$

Choose  $a$  from state  $s$  using a policy derived from  $q$ , such as  $\epsilon$ -greedy or softmax.

Choose  $a$  with prob.  $\frac{e^{q(s, a)}}{\sum_a e^{q(s, a)}}$

Repeat for each step in the episode:

- Take action  $a$ , observe  $r$  and  $s'$ .
- Choose  $a'$  using our  $q$  policy.
- $q(s, a) \leftarrow q(s, a) + \alpha(r + \gamma q(s, a') - q(s, a))$
- $a \leftarrow a', s \leftarrow s'$

if many  $u$ 's maximize  $q$ , choose among them with equal prob.  
With prob. select uniformly randomly from the full action set  $\mathcal{A}$ .

with prob  $1 - \epsilon$  choose an action  $a \in \arg \max_u q(s, u)$ ;

if many  $u$ 's maximize  $q$ , choose among them with equal prob.  
With prob. select uniformly randomly from the full action set  $\mathcal{A}$ .

Name Sarsa from  $(s, a, r, s', a')$

Func. Approx. Form:

Init:  $w \leftarrow 0$  (or maybe random)

Update:  $w \leftarrow w + \alpha(r + \gamma q_w(s', a') - q_w(s, a)) \cdot$

$$\frac{\partial q_w(s, a)}{\partial w}$$

## Sarsa properties:

- Converges a.s. to the optimal action value function if:

1) Tabular

2) Every  $(s,a)$  pair visited infinitely often

3) Adjust hyperparameters to move toward a greedy policy ( $\epsilon \rightarrow 0, \delta \rightarrow \text{large}$ )

) Fight each other.

Can use  $\epsilon_t = 1/t \dots$

But in practice  $\epsilon$  or  $\delta$  not adjusted.

"Greedy in the limit with infinite exploration" = GLIE

- What if  $\epsilon = 0$ ? Won't necessarily see all  $(s,a)$  pairs.

- What about a pessimistic initial value function?  
(initial  $q$  less than actual  $q$ ) Can easily get stuck on a value that increased a bit.

- What about an optimistic initial value function?  
Encourages exploration. Can help even if  $\epsilon > 0$ .

- Sarsa is "on-policy": Always estimating  $q$  for the current policy (the one generating actions).

- Linear func. approx:  
(converges a.s. (not clear to what)).

- Non-linear (general) func. approx: Can diverge.

Q-learning: off-policy TD control  $(s, a, r, s')$

Use this update:

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

$$w \leftarrow w + \alpha (r + \gamma \max_a q_w(s', a') - q_w(s, a)) \frac{\partial q_w(s, a)}{\partial w}$$