

AML, 5 Nov 2019

Monte Carlo methods to evaluate V_{π}

→ Temporal-Difference methods (TD)

Monte Carlo

Generate an episode w.r.t π

Calculate state values for the episode, right to left

Add new values to list $V(s)$ for each s

Finally take mean of each $V(s)$ list — initial estimates of $V(s)$ are discarded

TD:

"Bootstrapping" - like DP - *requires full model*
Uses current estimates of $V(s_t)$ to update $V(s_{t+1})$

TD - hybrid of DP & MC

If we were to do MC incrementally

$$V(s_t) \leftarrow V(s_t) + \alpha \left[\underbrace{G_t - V(s_t)}_{\text{observed error / deviation}} \right]$$

↑
learning rate

global value calculated for entire episode

Replace G_t by $\frac{R_{t+1} + \gamma V(s_{t+1})}{\text{observed reward} \quad \text{old estimate}}$

TD update rule

$$V(s_t) \leftarrow V(s_t) + \alpha \left[\frac{(R_{t+1} + \gamma V(s_{t+1})) - V(s_t)}{\text{error}} \right]$$

learning rate

TD [0] - zero lookahead

Tabular TD(0) for estimating v_π

Input: the policy π to be evaluated

Algorithm parameter: step size $\alpha \in (0, 1]$

Initialize $V(s)$, for all $s \in \mathcal{S}^+$, arbitrarily except that $V(\text{terminal}) = 0$

Loop for each episode:

 Initialize S

 Loop for each step of episode:

$A \leftarrow$ action given by π for S

 Take action A , observe R, S'

$V(S) \leftarrow V(S) + \alpha [R + \gamma V(S') - V(S)]$

$S \leftarrow S'$

 until S is terminal

$$v_{\pi}(s) = E_{\pi} [G_t | s_t = s]$$

Monte Carlo

$$R_{t+1} + \gamma G_{t+1} \\ \underline{\underline{v_{\pi}(s_{t+1})}}$$

$$E_{\pi} [R_{t+1} + \gamma v_{\pi}(s_{t+1}) | s_t = s] \quad \text{TD}[0]$$

TD-error $\delta_t \stackrel{\text{def}}{=} (R_{t+1} + \gamma V(s_{t+1})) - V(s_t)$

MC-error $G_t - V(s_t) = R_{t+1} + \gamma G_{t+1} - V(s_t) + \gamma V(s_{t+1}) - \gamma V(s_{t+1})$


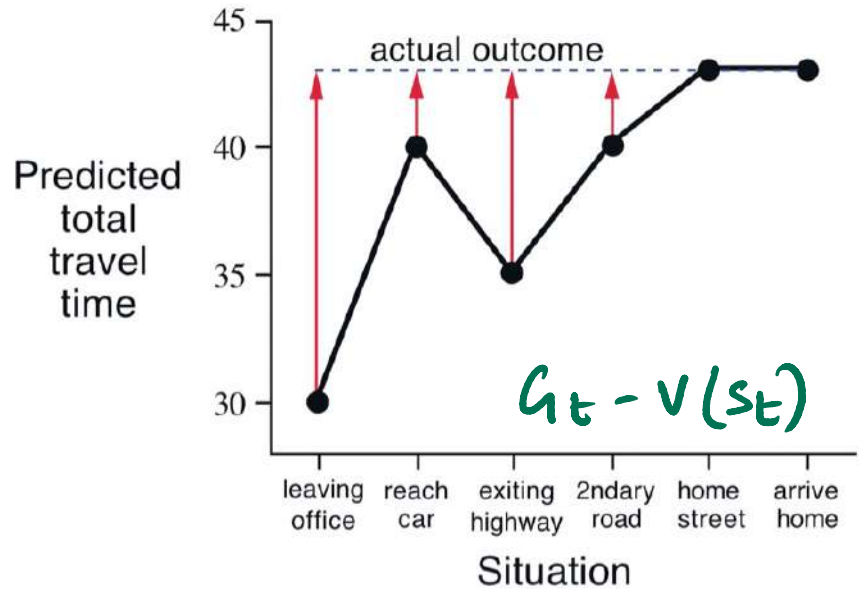
MC-error $G_t - V(S_t) = \underbrace{R_{t+1}} + \underbrace{rG_{t+1}} - \underbrace{V(S_t)} + \underbrace{rV(S_{t+1})} - \underbrace{rV(S_{t+1})}$

$$= \frac{(R_{t+1} + rV(S_{t+1}) - V(S_t))}{\delta_t} + \frac{r(G_{t+1} - V(S_{t+1}))}{\delta_t}$$

$$= \delta_t + \gamma \cdot \delta_{t+1} + \gamma^2 \delta_{t+2} \dots$$

$$= \sum_{k=t}^{T-1} \gamma^{k-t} \delta_k$$

<i>State</i>	<i>Elapsed Time (minutes)</i>	<i>Predicted Time to Go</i>	<i>Predicted Total Time</i>
leaving office, friday at 6	0	30	30
reach car, raining	5	35	40
exiting highway	20	15	35
2ndary road, behind truck	30	10	40
entering home street	40	3	43
arrive home	43	0	43




Advantage of TD vs MC

- Incremental update of v
- No need to wait till episode ends
 - May not even have finite episodes

Both TD & MC converge asymptotically to correct values

Which is faster?

- Not been proven
- In practice TD appears faster



$1/6$


$2/6$

$V = 1/2$

$4/6$

$5/6$

$TD[0]$
starting
with $V=0.5$
everywhere




theoretical
value

True
values

State

Empirical RMS error, averaged over states



Batch learning

Consider entire set of episodes as a single batch update

MC & TD give different answers


Example 8 observations

A, 0, B, 0
B, 1
B, 1
B, 1

B, 1
B, 1
B, 1
B, 0

$$V(B) = \frac{6}{8} = \frac{3}{4}$$

What is $V(A)$? $\left\{ \begin{array}{l} \text{MC sets } V(A) = 0 \\ \text{TD sets } V(A) = V(B) \end{array} \right.$



MDP learned
by TD[0]

MC — Best estimate wrt Mean Square Error
of observations

TD[0] — Best estimate wrt. MLE

From value estimation to policy iteration

- MC : On policy - Same π generates runs & gets update
Off policy - runs are generated independent
of π to be updated

On policy TD

Typically switch from estimating $V_{\pi}(s)$ to $q_{\pi}(s,a)$

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha \left[R_{t+1} + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t) \right]$$

Update is a function of $(s_t, a_t, R_{t+1}, s_{t+1}, a_{t+1}) \Rightarrow$ SARSA

Sarsa (on-policy TD control) for estimating $Q \approx q_*$

Algorithm parameters: step size $\alpha \in (0, 1]$, small $\varepsilon > 0$

Initialize $Q(s, a)$, for all $s \in \mathcal{S}^+$, $a \in \mathcal{A}(s)$, arbitrarily except that $Q(\text{terminal}, \cdot) = 0$

Loop for each episode:

Initialize S

Choose A from S using policy derived from Q (e.g., ε -greedy) $Q \rightarrow \pi$

Loop for each step of episode:

Take action A , observe R, S' *on policy*

Choose A' from S' using policy derived from Q (e.g., ε -greedy) *Update π*

$Q(S, A) \leftarrow Q(S, A) + \alpha [R + \gamma Q(S', A') - Q(S, A)]$ *Update Q*

$S \leftarrow S'; A \leftarrow A';$

until S is terminal

Update $Q(S_t, A_t) \longrightarrow$ Update π
←

Off Policy

Q-learning

$$Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \alpha \left[R_{t+1} + \gamma \max_a Q(S_{t+1}, A_{t+1}) - Q(S_t, A_t) \right]$$

next action is best
wrt Q, not
determined by π

Q-learning (off-policy TD control) for estimating $\pi \approx \pi_*$

Algorithm parameters: step size $\alpha \in (0, 1]$, small $\varepsilon > 0$

Initialize $Q(s, a)$, for all $s \in \mathcal{S}^+$, $a \in \mathcal{A}(s)$, arbitrarily except that $Q(\text{terminal}, \cdot) = 0$

Loop for each episode:

 Initialize S

 Loop for each step of episode:

 Choose A from S using policy derived from Q (e.g., ε -greedy)

 Take action A , observe R, S'

$$Q(S, A) \leftarrow Q(S, A) + \alpha [R + \gamma \max_a Q(S', a) - Q(S, A)]$$

$S \leftarrow S'$

 until S is terminal


Variante η SARSA - Expected SARSA

$$Q(s_t, A_t) \leftarrow Q(s_t, A_t) + \alpha \left[R_{t+1} + \gamma E_{\pi} \left[Q(s_{t+1}, A_{t+1}) \mid s_{t+1} \right] - Q(s_t, A_t) \right]$$

$$\sum_a \pi(a \mid s_{t+1}) Q(s_{t+1}, a)$$

Maximization Bias

-tends to produce inflated estimates



Double learning - Simultaneously create 2 estimates,
randomly use one to update other.

Double Q-learning, for estimating $Q_1 \approx Q_2 \approx q_*$

Algorithm parameters: step size $\alpha \in (0, 1]$, small $\varepsilon > 0$

Initialize $Q_1(s, a)$ and $Q_2(s, a)$, for all $s \in \mathcal{S}^+, a \in \mathcal{A}(s)$, such that $Q(\text{terminal}, \cdot) = 0$

Loop for each episode:

Initialize S

Loop for each step of episode:

Choose A from S using the policy ε -greedy in $Q_1 + Q_2$

Take action A , observe R, S'

With 0.5 probability:

$$Q_1(S, A) \leftarrow Q_1(S, A) + \alpha \left(R + \gamma Q_2(S', \arg \max_a Q_1(S', a)) - Q_1(S, A) \right)$$

else:

$$Q_2(S, A) \leftarrow Q_2(S, A) + \alpha \left(R + \gamma Q_1(S', \arg \max_a Q_2(S', a)) - Q_2(S, A) \right)$$

$S \leftarrow S'$

until S is terminal