

Markov Decision Process, Optimal Solutions, Monte Carlo Methods

Milan Straka

 October 11, 2021



EUROPEAN UNION
European Structural and Investment Fund
Operational Programme Research,
Development and Education

Charles University in Prague
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated




Figure 3.1 of "Reinforcement Learning: An Introduction, Second Edition".

A **Markov decision process** (MDP) is a quadruple $(\mathcal{S}, \mathcal{A}, p, \gamma)$, where:

- \mathcal{S} is a set of states,
- \mathcal{A} is a set of actions,
- $p(S_{t+1} = s', R_{t+1} = r | S_t = s, A_t = a)$ is a probability that action $a \in \mathcal{A}$ will lead from state $s \in \mathcal{S}$ to $s' \in \mathcal{S}$, producing a **reward** $r \in \mathbb{R}$,
- $\gamma \in [0, 1]$ is a **discount factor**.

Let a **return** G_t be $G_t \stackrel{\text{def}}{=} \sum_{k=0}^{\infty} \gamma^k R_{t+1+k}$. The goal is to optimize $\mathbb{E}[G_0]$.

Partially Observable MDPs




Figure 3.1 of "Reinforcement Learning: An Introduction, Second Edition".

Partially observable Markov decision process extends the Markov decision process to a sextuple $(\mathcal{S}, \mathcal{A}, p, \gamma, \mathcal{O}, o)$, where in addition to an MDP,

- \mathcal{O} is a set of observations,
- $o(O_t | S_t, A_{t-1})$ is an observation model, which is used as agent input instead of S_t .

Multi-armed Bandits as MDP

To formulate n -armed bandits problem as MDP, we do not need states. Therefore, we could formulate it as:

- one-element set of states, $\mathcal{S} = \{S\}$;
- an action for every arm, $\mathcal{A} = \{a_1, a_2, \dots, a_n\}$;
- assuming every arm produces rewards with a distribution of $\mathcal{N}(\mu_i, \sigma_i^2)$, the MDP dynamics function p is defined as

$$p(S, r | S, a_i) = \mathcal{N}(r | \mu_i, \sigma_i^2).$$

One possibility to introduce states in multi-armed bandits problem is to consider a separate reward distribution for every state. Such generalization is called **Contextualized Bandits** problem. Assuming state transitions are independent on rewards and given by a distribution $next(s)$, the MDP dynamics function for contextualized bandits problem is given by

$$p(s', r | s, a_i) = \mathcal{N}(r | \mu_{i,s}, \sigma_{i,s}^2) \cdot next(s' | s).$$

Episodic and Continuing Tasks

If the agent-environment interaction naturally breaks into independent subsequences, usually called **episodes**, we talk about **episodic tasks**. Each episode then ends in a special **terminal state**, followed by a reset to a starting state (either always the same, or sampled from a distribution of starting states).

In episodic tasks, it is often the case that every episode ends in at most H steps. These **finite-horizon tasks** then can use discount factor $\gamma = 1$, because the return $G \stackrel{\text{def}}{=} \sum_{t=0}^H \gamma^t R_{t+1}$ is well defined.

If the agent-environment interaction goes on and on without a limit, we instead talk about **continuing tasks**. In this case, the discount factor γ needs to be sharply smaller than 1.

(State-)Value and Action-Value Functions

A **policy** π computes a distribution of actions in a given state, i.e., $\pi(a|s)$ corresponds to a probability of performing an action a in state s .

To evaluate a quality of a policy, we define **value function** $v_\pi(s)$, or **state-value function**, as

$$\begin{aligned} v_\pi(s) &\stackrel{\text{def}}{=} \mathbb{E}_\pi [G_t | S_t = s] = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \mid S_t = s \right] \\ &= \mathbb{E}_{A_t \sim \pi(s)} \mathbb{E}_{S_{t+1}, R_{t+1} \sim p(s, A_t)} \left[R_{t+1} + \gamma \mathbb{E}_{A_{t+1} \sim \pi(S_{t+1})} \mathbb{E}_{S_{t+2}, R_{t+2} \sim p(S_{t+1}, A_{t+1})} [R_{t+2} + \dots] \right] \end{aligned}$$

An **action-value function** for a policy π is defined analogously as

$$q_\pi(s, a) \stackrel{\text{def}}{=} \mathbb{E}_\pi [G_t | S_t = s, A_t = a] = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k R_{t+k+1} \mid S_t = s, A_t = a \right].$$

The value function and state-value function can be of course expressed using one another:

$$v_\pi(s) = \mathbb{E}_{a \sim \pi} [q_\pi(s, a)], \quad q_\pi(s, a) = \mathbb{E}_{s', r \sim p} [r + \gamma v_\pi(s')].$$

Optimal state-value function is defined as

$$v_*(s) \stackrel{\text{def}}{=} \max_{\pi} v_{\pi}(s),$$

analogously

$$q_*(s, a) \stackrel{\text{def}}{=} \max_{\pi} q_{\pi}(s, a).$$

Any policy π_* with $v_{\pi_*} = v_*$ is called an **optimal policy**. Such policy can be defined as $\pi_*(s) \stackrel{\text{def}}{=} \arg \max_a q_*(s, a) = \arg \max_a \mathbb{E}[R_{t+1} + \gamma v_*(S_{t+1}) | S_t = s, A_t = a]$. When multiple actions maximize $q_*(s, a)$, the optimal policy can stochastically choose any of them.

Existence

In finite-horizon tasks or if $\gamma < 1$, there always exists a unique optimal state-value function, a unique optimal action-value function, and a (not necessarily unique) optimal policy.

Dynamic programming is an approach devised by Richard Bellman in 1950s.

To apply it to MDP, we now consider finite-horizon problems with finite number of states \mathcal{S} , finite number of actions \mathcal{A} , and known MDP dynamics p . Note that without loss of generality, we can assume that every episode takes exactly H steps (by introducing a suitable absorbing state, if necessary).

The following recursion is usually called the *Bellman equation*:

$$\begin{aligned} v_*(s) &= \max_a \mathbb{E}[R_{t+1} + \gamma v_*(S_{t+1}) | S_t = s, A_t = a] \\ &= \max_a \sum_{s', r} p(s', r | s, a) [r + \gamma v_*(s')]. \end{aligned}$$

It must hold for an optimal value function in a MDP, because future decisions do not depend on the current one. Therefore, the optimal policy can be expressed as one action followed by optimal policy from the resulting state.

To turn the Bellman equation into an algorithm, we change the equal signs to assignments:


$$v_0(s) \leftarrow \begin{cases} 0 & \text{for the terminal state } s \\ -\infty & \text{otherwise} \end{cases}$$
$$v_{k+1}(s) \leftarrow \max_a \mathbb{E} [R_{t+1} + \gamma v_k(S_{t+1}) | S_t = s, A_t = a].$$

In a finite-horizon task with H steps, the optimal value function is reached after H iterations of the above assignment:

- We can show by induction that $v_k(s)$ is the maximum return reachable from state s in last k steps of an episode.
- If every episode ends in at most H steps, then v_{H+1} must be equal to v_H .

Relations to Graph Algorithms

Searching for optimal value functions of deterministic problems is in fact search for the shortest path in a suitable graph.



https://upload.wikimedia.org/wikipedia/commons/a/a0/Convolutional_code_trellis_diagram.svg

Bellman-Ford-Moore Algorithm

$$v_{k+1}(s) \leftarrow \max_a \mathbb{E} [R_{t+1} + \gamma v_k(S_{t+1}) | S_t = s, A_t = a].$$

Bellman-Ford-Moore algorithm:

```
# input: graph `g`, initial vertex `s`
for v in g.vertices:
    d[v] = 0 if v == s else +∞

for iteration in range(len(g.vertices) - 1):
    for e in g.edges:
        if d[e.source] + e.length < d[e.target]:
            d[e.target] = d[e.source] + e.length
```

Uniqueness of Bellman Equation Solution

Not only does the optimal value function fulfill the Bellman equation in the current settings, the converse is also true: If a value function satisfies the Bellman equation, it is optimal.

To sketch the proof of the statement, consider for a contradiction that some solution of Bellman equation is not an optimal value function. Therefore, there exist states with different than optimal values.

Among those states, at least one is always the last one on a trajectory (if it is present). However, it can be easily verified that if its value is not optimal, Bellman equation cannot hold in this state.

Bellman Backup Operator

Our goal is now to handle also infinite-horizon tasks, using discount factor of $\gamma < 1$. However, we still assume finite number of states and actions.

For any value function $v \in \mathbb{R}^{|\mathcal{S}|}$ we define *Bellman backup operator* $B : \mathbb{R}^{|\mathcal{S}|} \rightarrow \mathbb{R}^{|\mathcal{S}|}$ as

$$Bv(s) \stackrel{\text{def}}{=} \max_a \mathbb{E} [R_{t+1} + \gamma v(S_{t+1}) | S_t = s, A_t = a].$$

Considering the supremum norm $\|x\|_\infty \stackrel{\text{def}}{=} \sup_s |x(s)|$, we will show that Bellman backup operator is a *contraction* (even for infinite number of states), i.e.,

$$\sup_s |Bv_1(s) - Bv_2(s)| = \|Bv_1 - Bv_2\|_\infty \leq \gamma \|v_1 - v_2\|_\infty.$$

Applying the Banach fixed-point theorem on the normed vector space $\mathbb{R}^{|\mathcal{S}|}$ with the supremum norm then yields that there exists a *unique value function* v_* such that $Bv_* = v_*$.

Such a unique v_* is the *optimal value function*, because it satisfies the Bellman equation.

Bellman Backup Operator

Furthermore, iterative application of B on arbitrary v converges to v_* , because

$$\|Bv - v_*\|_\infty = \|Bv - Bv_*\|_\infty \leq \gamma \|v - v_*\|,$$

and therefore $B^n v \rightarrow v_*$.

Value Iteration Algorithm

We can turn the iterative application of Bellman backup operator into an algorithm.

$$Bv(s) \stackrel{\text{def}}{=} \max_a \mathbb{E} [R_{t+1} + \gamma v(S_{t+1}) | S_t = s, A_t = a]$$

Value Iteration, for estimating $\pi \approx \pi_*$

Algorithm parameter: a small threshold $\theta > 0$ determining accuracy of estimation
 Initialize $V(s)$, for all $s \in \mathcal{S}$, arbitrarily except that $V(\text{terminal}) = 0$

Loop:

```

|  $\Delta \leftarrow 0$ 
| Loop for each  $s \in \mathcal{S}$ :
|    $v \leftarrow V(s)$ 
|    $V(s) \leftarrow \max_a \sum_{s',r} p(s', r | s, a) [r + \gamma V(s')]$ 
|    $\Delta \leftarrow \max(\Delta, |v - V(s)|)$ 
until  $\Delta < \theta$ 
  
```

Output a deterministic policy, $\pi \approx \pi_*$, such that
 $\pi(s) = \arg \max_a \sum_{s',r} p(s', r | s, a) [r + \gamma V(s')]$

Modification of Algorithm 4.4 of "Reinforcement Learning: An Introduction, Second Edition" (replacing S+ by S).

Although we have described the so-called *synchronous* implementation requiring two arrays for v and Bv , usual implementations are *asynchronous* and modify the value function in place (if a fixed ordering is used, usually such value iteration is called *Gauss-Seidel*).

- for $s \in \mathcal{S}$ in some fixed order:
 - $v(s) \leftarrow \max_a \mathbb{E}[R_{t+1} + \gamma v(S_{t+1}) | S_t = s, A_t = a]$

Even with such asynchronous update, value iteration can be proven to converge, and usually performs better in practice.

For example, the Bellman-Ford-Moore algorithm also updates the distances in-place. In the case of dynamic programming, we can extend the invariant from “ $v_k(s)$ is the maximum return reachable from state s in last k steps of an episode” to include not only all trajectories of k steps, but also any number of longer trajectories.

If you are interested, try proving that the above Gauss-Seidel iteration is also a contraction.

Bellman Backup Operator as a Contraction

To show that Bellman backup operator is a contraction, we proceed as follows:

$$\begin{aligned}
 \|Bv_1 - Bv_2\|_\infty &= \left\| \max_a \mathbb{E}[R_{t+1} + \gamma v_1(S_{t+1})] - \max_a \mathbb{E}[R_{t+1} + \gamma v_2(S_{t+1})] \right\|_\infty \\
 &= \left\| \max_a \left(\mathbb{E}[R_{t+1} + \gamma v_1(S_{t+1})] - \max_a \mathbb{E}[R_{t+1} + \gamma v_2(S_{t+1})] \right) \right\|_\infty \\
 &\leq \left\| \max_a \left(\mathbb{E}[R_{t+1} + \gamma v_1(S_{t+1})] - \mathbb{E}[R_{t+1} + \gamma v_2(S_{t+1})] \right) \right\|_\infty \\
 &= \max_a \left(\left\| \mathbb{E}[R_{t+1} + \gamma v_1(S_{t+1})] - \mathbb{E}[R_{t+1} + \gamma v_2(S_{t+1})] \right\|_\infty \right) \\
 &= \max_a \left(\left\| \sum_{s',r} p(s',r|s,a) \gamma (v_1(s') - v_2(s')) \right\|_\infty \right) \\
 &= \gamma \max_a \left(\left\| \sum_{s'} p(s'|s,a) (v_1(s') - v_2(s')) \right\|_\infty \right) \\
 &\leq \gamma \|v_1 - v_2\|
 \end{aligned}$$

where the last line follows from the fact that for any s and a , $\sum_{s'} p(s'|s,a)$ sums to 1.

Speed of Convergence

Assuming maximum reward is R_{\max} , we have that

$$v_*(s) \leq \sum_{t=0}^{\infty} \gamma^t R_{\max} = \frac{R_{\max}}{1 - \gamma}.$$

Starting with $v(s) \leftarrow 0$, we have

$$\|B^k v - v_*\|_{\infty} \leq \gamma^k \|v - v_*\|_{\infty} \leq \gamma^k \frac{R_{\max}}{1 - \gamma}.$$

Compare to finite-horizon case, where $B^T v = v_*$.

Value Iteration Example

Consider a simple betting game, where a gambler repeatedly bets on the outcome of a coin flip (with a given win probability), either losing their stake or winning the same amount of coins that was bet. The gambler wins if they obtain 100 coins, and lose if they run out of money.

We can formulate the problem as an undiscounted episodic MDP. The states are the coins owned by the gambler, $\{1, \dots, 99\}$, and actions are the stakes $\{1, \dots, \min(s, 100 - s)\}$.

The reward is $+1$ when reaching 100 and 0 otherwise.

The state-value function then gives probability of winning from each state, and policy prescribes a stake with a given capital.

Value Iteration Example

For a coin flip win probability 40%, the value iteration proceeds as follows.




Figure 4.3 of "Reinforcement Learning: An Introduction, Second Edition".

Policy Iteration Algorithm

We now propose another approach of computing optimal policy. The approach, called **policy iteration**, consists of repeatedly performing policy **evaluation** and policy **improvement**.

Policy Evaluation

Given a policy π , policy evaluation computes v_π .

Recall that

$$\begin{aligned} v_\pi(s) &\stackrel{\text{def}}{=} \mathbb{E}_\pi [G_t | S_t = s] \\ &= \mathbb{E}_\pi [R_{t+1} + \gamma v_\pi(S_{t+1}) | S_t = s] \\ &= \sum_a \pi(a|s) \sum_{s', r} p(s', r | s, a) [r + \gamma v_\pi(s')]. \end{aligned}$$

If the dynamics of the MDP p is known, the above is a system of linear equations, and therefore, v_π can be computed exactly.

The equation

$$v_{\pi}(s) = \sum_a \pi(a|s) \sum_{s',r} p(s', r|s, a) [r + \gamma v_{\pi}(s')]$$

is called *Bellman equation for v_{π}* and analogously to Bellman optimality equation, it can be proven that

- under the same assumptions as before ($\gamma < 1$ or termination), v_{π} exists and is unique;
- v_{π} is a fixed point of the Bellman equation

$$v_{k+1}(s) = \sum_a \pi(a|s) \sum_{s',r} p(s', r|s, a) [r + \gamma v_k(s')];$$

- iterative application of the Bellman equation to any v converges to v_{π} (the proof is easier than for the optimality equation, because v_{π} is defined using an expectation and expectations are linear, so we get the first half of the proof “for free”).

Iterative Policy Evaluation, for estimating $V \approx v_\pi$

Input π , the policy to be evaluated

Algorithm parameter: a small threshold $\theta > 0$ determining accuracy of estimation

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

Loop:

$\Delta \leftarrow 0$

Loop for each $s \in \mathcal{S}$:

$v \leftarrow V(s)$

$V(s) \leftarrow \sum_a \pi(a|s) \sum_{s',r} p(s',r|s,a) [r + \gamma V(s')]$

$\Delta \leftarrow \max(\Delta, |v - V(s)|)$

until $\Delta < \theta$

Modification of Algorithm 4.1 of "Reinforcement Learning: An Introduction, Second Edition" (replacing S+ by S).

Given π and computed v_π , we would like to **improve** the policy. A straightforward way to do so is to define a policy using a *greedy* action

$$\begin{aligned}\pi'(s) &\stackrel{\text{def}}{=} \arg \max_a q_\pi(s, a) \\ &= \arg \max_a \sum_{s', r} p(s', r | s, a) [r + \gamma v_\pi(s')].\end{aligned}$$

For such π' , by construction it obviously holds that

$$q_\pi(s, \pi'(s)) \geq v_\pi(s).$$

Policy Improvement Theorem


Let π and π' be any pair of deterministic policies, such that $q_\pi(s, \pi'(s)) \geq v_\pi(s)$.

Then for all states s , $v_{\pi'}(s) \geq v_\pi(s)$.

The proof is straightforward, we repeatedly expand q_π and use the assumption of the policy improvement theorem:

$$\begin{aligned}
 v_\pi(s) &\leq q_\pi(s, \pi'(s)) \\
 &= \mathbb{E}[R_{t+1} + \gamma v_\pi(S_{t+1}) | S_t = s, A_t = \pi'(s)] \\
 &= \mathbb{E}_{\pi'}[R_{t+1} + \gamma v_\pi(S_{t+1}) | S_t = s] \\
 &\leq \mathbb{E}_{\pi'}[R_{t+1} + \gamma q_\pi(S_{t+1}, \pi'(S_{t+1})) | S_t = s] \\
 &= \mathbb{E}_{\pi'}[R_{t+1} + \gamma \mathbb{E}[R_{t+2} + \gamma v_\pi(S_{t+2}) | S_{t+1}, A_{t+1} = \pi'(S_{t+1})] | S_t = s] \\
 &= \mathbb{E}_{\pi'}[R_{t+1} + \gamma R_{t+2} + \gamma^2 v_\pi(S_{t+2}) | S_t = s] \\
 &\dots \\
 &\leq \mathbb{E}_{\pi'}[R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots | S_t = s] = v_{\pi'}
 \end{aligned}$$

Policy Improvement Example



	1	2	3
4	5	6	7
8	9	10	11
12	13	14	

$R_t = -1$
on all transitions

Example 4.1 of "Reinforcement Learning: An Introduction, Second Edition".




Figure 4.1 of "Reinforcement Learning: An Introduction, Second Edition".

Policy Iteration Algorithm

Policy iteration consists of repeatedly performing policy evaluation and policy improvement:

$$\pi_0 \xrightarrow{E} v_{\pi_0} \xrightarrow{I} \pi_1 \xrightarrow{E} v_{\pi_1} \xrightarrow{I} \pi_2 \xrightarrow{E} v_{\pi_2} \xrightarrow{I} \dots \xrightarrow{I} \pi_* \xrightarrow{E} v_{\pi_*}.$$

The result is a sequence of monotonically improving policies π_i . Note that when $\pi' = \pi$, also $v_{\pi'} = v_{\pi}$, which means Bellman optimality equation is fulfilled and both v_{π} and π are optimal.

Considering that there is only a finite number of policies, the optimal policy and optimal value function can be computed in finite time (contrary to value iteration, where the convergence is only asymptotic).

Note that when evaluating policy π_{k+1} , we usually start with v_{π_k} , which is assumed to be a good approximation to $v_{\pi_{k+1}}$.

Policy Iteration (using iterative policy evaluation) for estimating $\pi \approx \pi_*$

1. Initialization

$V(s) \in \mathbb{R}$ and $\pi(s) \in \mathcal{A}(s)$ arbitrarily for all $s \in \mathcal{S}$

2. Policy Evaluation

Loop:

$\Delta \leftarrow 0$

Loop for each $s \in \mathcal{S}$:

$v \leftarrow V(s)$

$V(s) \leftarrow \sum_{s',r} p(s',r|s,\pi(s)) [r + \gamma V(s')]$

$\Delta \leftarrow \max(\Delta, |v - V(s)|)$

until $\Delta < \theta$ (a small positive number determining the accuracy of estimation)

3. Policy Improvement

policy-stable \leftarrow *true*

For each $s \in \mathcal{S}$:

old-action $\leftarrow \pi(s)$

$\pi(s) \leftarrow \operatorname{argmax}_a \sum_{s',r} p(s',r|s,a) [r + \gamma V(s')]$

If *old-action* $\neq \pi(s)$, then *policy-stable* \leftarrow *false*

If *policy-stable*, then stop and return $V \approx v_*$ and $\pi \approx \pi_*$; else go to 2

Algorithm 4.3 of "Reinforcement Learning: An Introduction, Second Edition".

Note that value iteration is in fact a policy iteration, where policy evaluation is performed only for one step:

$$\pi'(s) = \arg \max_a \sum_{s',r} p(s', r | s, a) [r + \gamma v(s')] \quad (\textit{policy improvement})$$

$$v'(s) = \sum_a \pi'(a|s) \sum_{s',r} p(s', r | s, a) [r + \gamma v(s')] \quad (\textit{one step of policy evaluation})$$

Substituting the former into the latter, we get

$$v'(s) = \max_a \sum_{s',r} p(s', r | s, a) [r + \gamma v(s')] = Bv(s).$$

Generalized Policy Iteration

Therefore, it seems that to achieve convergence, it is not necessary to perform the policy evaluation exactly.

Generalized Policy Evaluation is a general idea of interleaving policy evaluation and policy improvement at various granularity.




Figure in Section 4.6 of "Reinforcement Learning: An Introduction, Second Edition".




Figure in Section 4.6 of "Reinforcement Learning: An Introduction, Second Edition".

If both processes stabilize, we know we have obtained optimal policy.

Monte Carlo methods are based on estimating returns from complete episodes. Furthermore, if the model (of the environment) is not known, we need to estimate returns for the action-value function q instead of v .

We can formulate Monte Carlo methods in the generalized policy improvement framework. Keeping estimated returns for the action-value function, we perform policy evaluation by sampling one episode according to current policy. We then update the action-value function by averaging over the observed returns, including the currently sampled episode.

To hope for convergence, we need to visit each state infinitely many times. One of the simplest way to achieve that is to assume *exploring starts*, where we randomly select the first state and first action, each pair with nonzero probability.

Furthermore, if a state-action pair appears multiple times in one episode, the sampled returns are not independent. The literature distinguishes two cases:

- *first visit*: only the first occurrence of a state-action pair in an episode is considered
- *every visit*: all occurrences of a state-action pair are considered.

Even though first-visit is easier to analyze, it can be proven that for both approaches, policy evaluation converges. Contrary to the Reinforcement Learning: An Introduction book, which presents first-visit algorithms, we use every-visit.

Monte Carlo ES (Exploring Starts), for estimating $\pi \approx \pi_*$

Initialize:

$\pi(s) \in \mathcal{A}(s)$ (arbitrarily), for all $s \in \mathcal{S}$

$Q(s, a) \in \mathbb{R}$ (arbitrarily), for all $s \in \mathcal{S}$, $a \in \mathcal{A}(s)$

$Returns(s, a) \leftarrow$ empty list, for all $s \in \mathcal{S}$, $a \in \mathcal{A}(s)$

Loop forever (for each episode):

Choose $S_0 \in \mathcal{S}$, $A_0 \in \mathcal{A}(S_0)$ randomly such that all pairs have probability > 0

Generate an episode from S_0, A_0 , following π : $S_0, A_0, R_1, \dots, S_{T-1}, A_{T-1}, R_T$

$G \leftarrow 0$

Loop for each step of episode, $t = T-1, T-2, \dots, 0$:

$G \leftarrow \gamma G + R_{t+1}$

Append G to $Returns(S_t, A_t)$

$Q(S_t, A_t) \leftarrow \text{average}(Returns(S_t, A_t))$

$\pi(S_t) \leftarrow \text{argmax}_a Q(S_t, a)$

Modification of algorithm 5.3 of "Reinforcement Learning: An Introduction, Second Edition" from first-visit to every-visit.

Monte Carlo and ε -soft Policies

The problem with exploring starts is that in many situations, we either cannot start in an arbitrary state, or it is impractical.

A policy is called ε -soft, if

$$\pi(a|s) \geq \frac{\varepsilon}{|\mathcal{A}(s)|}.$$

and we call it ε -greedy, if one action has a maximum probability of $1 - \varepsilon + \frac{\varepsilon}{|\mathcal{A}(s)|}$.

The policy improvement theorem can be proved also for the class of ε -soft policies, and using ε -greedy policy in policy improvement step, policy iteration has the same convergence properties. (We can embed the ε -soft behaviour “inside” the environment and prove equivalence.)

On-policy every-visit Monte Carlo for ε -soft Policies

Algorithm parameter: small $\varepsilon > 0$

Initialize $Q(s, a) \in \mathbb{R}$ arbitrarily (usually to 0), for all $s \in \mathcal{S}, a \in \mathcal{A}$

Initialize $C(s, a) \in \mathbb{Z}$ to 0, for all $s \in \mathcal{S}, a \in \mathcal{A}$

Repeat forever (for each episode):

- Generate an episode $S_0, A_0, R_1, \dots, S_{T-1}, A_{T-1}, R_T$, by generating actions as follows:
 - With probability ε , generate a random uniform action
 - Otherwise, set $A_t \stackrel{\text{def}}{=} \arg \max_a Q(S_t, a)$
- $G \leftarrow 0$
- For each $t = T - 1, T - 2, \dots, 0$:
 - $G \leftarrow \gamma G + R_{t+1}$
 - $C(S_t, A_t) \leftarrow C(S_t, A_t) + 1$
 - $Q(S_t, A_t) \leftarrow Q(S_t, A_t) + \frac{1}{C(S_t, A_t)} (G - Q(S_t, A_t))$