

Author (all other notes): Nikhil Sharma

Author (Bayes' Nets notes): Josh Hug and Jacky Liang, edited by Regina Wang

Author (Logic notes): Henry Zhu, edited by Peyrin Kao

Credit (Machine Learning and Logic notes): Some sections adapted from the textbook *Artificial Intelligence: A Modern Approach*.

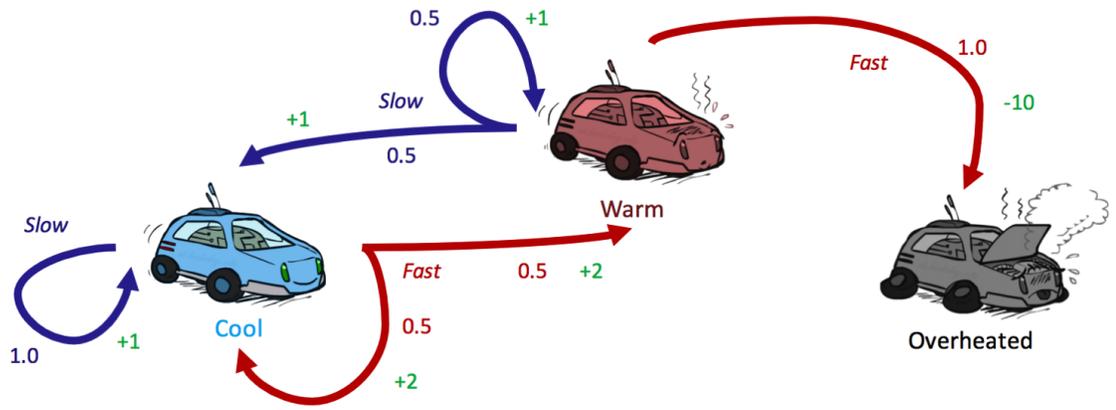
Last updated: March 8, 2024

Markov Decision Processes

A Markov Decision Process is defined by several properties:

- A set of states S . States in MDPs are represented in the same way as states in traditional search problems.
- A set of actions A . Actions in MDPs are also represented in the same way as in traditional search problems.
- A start state.
- Possibly one or more terminal states.
- Possibly a **discount factor** γ . We'll cover discount factors shortly.
- A **transition function** $T(s, a, s')$. Since we have introduced the possibility of nondeterministic actions, we need a way to delineate the likelihood of the possible outcomes after taking any given action from any given state. The transition function for a MDP does exactly this - it's a probability function which represents the probability that an agent taking an action $a \in A$ from a state $s \in S$ ends up in a state $s' \in S$.
- A **reward function** $R(s, a, s')$. Typically, MDPs are modeled with small "living" rewards at each step to reward an agent's survival, along with large rewards for arriving at a terminal state. Rewards may be positive or negative depending on whether or not they benefit the agent in question, and the agent's objective is naturally to acquire the maximum reward possible before arriving at some terminal state.

Constructing a MDP for a situation is quite similar to constructing a state-space graph for a search problem, with a couple additional caveats. Consider the motivating example of a racecar:



There are three possible states, $S = \{cool, warm, overheated\}$, and two possible actions $A = \{slow, fast\}$. Just like in a state-space graph, each of the three states is represented by a node, with edges representing actions. *Overheated* is a terminal state, since once a racecar agent arrives at this state, it can no longer perform any actions for further rewards (it's a *sink state* in the MDP and has no outgoing edges). Notably, for nondeterministic actions, there are multiple edges representing the same action from the same state with differing successor states. Each edge is annotated not only with the action it represents, but also a transition probability and corresponding reward. These are summarized below:

- **Transition Function:** $T(s, a, s')$
 - $T(cool, slow, cool) = 1$
 - $T(warm, slow, cool) = 0.5$
 - $T(warm, slow, warm) = 0.5$
 - $T(cool, fast, cool) = 0.5$
 - $T(cool, fast, warm) = 0.5$
 - $T(warm, fast, overheated) = 1$
- **Reward Function:** $R(s, a, s')$
 - $R(cool, slow, cool) = 1$
 - $R(warm, slow, cool) = 1$
 - $R(warm, slow, warm) = 1$
 - $R(cool, fast, cool) = 2$
 - $R(cool, fast, warm) = 2$
 - $R(warm, fast, overheated) = -10$

We represent the movement of an agent through different MDP states over time with discrete **timesteps**, defining $s_t \in S$ and $a_t \in A$ as the state in which an agent exists and the action which an agent takes at timestep t respectively. An agent starts in state s_0 at timestep 0, and takes an action at every timestep. The movement of an agent through a MDP can thus be modeled as follows:

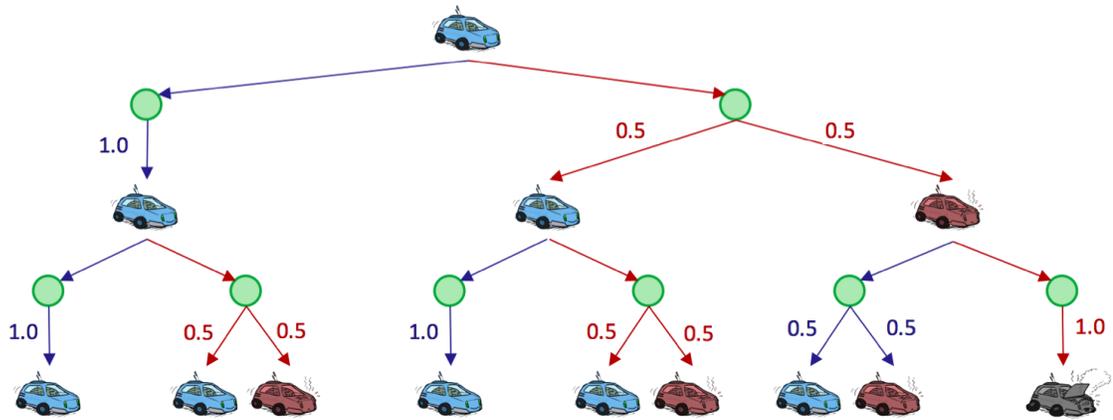
$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \xrightarrow{a_2} s_3 \xrightarrow{a_3} \dots$$

Additionally, knowing that an agent's goal is to maximize it's reward across all timesteps, we can correspondingly express this mathematically as a maximization of the following utility function:

$$U([s_0, a_0, s_1, a_1, s_2, \dots]) = R(s_0, a_0, s_1) + R(s_1, a_1, s_2) + R(s_2, a_2, s_3) + \dots$$

Markov decision processes, like state-space graphs, can be unraveled into search trees. Uncertainty is modeled in these search trees with **Q-states**, also known as **action states**, essentially identical to expectimax chance nodes. This is a fitting choice, as Q-states use probabilities to model the uncertainty that the environment will land an agent in a given state just as expectimax chance nodes use probabilities to model the uncertainty that adversarial agents will land our agent in a given state through the move these agents select. The Qstate represented by having taken action a from state s is notated as the tuple (s, a) .

Observe the unraveled search tree for our racecar, truncated to depth-2:



The green nodes represent Q-states, where an action has been taken from a state but has yet to be resolved into a successor state. It's important to understand that agents spend zero timesteps in Q-states, and that they are simply a construct created for ease of representation and development of MDP algorithms.

Finite Horizons and Discounting

There is an inherent problem with our racecar MDP - we haven't placed any time constraints on the number of timesteps for which a racecar can take actions and collect rewards. With our current formulation, it could routinely choose $a = \text{slow}$ at every timestep forever, safely and effectively obtaining infinite reward without any risk of overheating. This is prevented by the introduction of **finite horizons** and/or **discount factors**. An MDP enforcing a finite horizon is simple - it essentially defines a "lifetime" for agents, which gives them some set number of timesteps n to accrue as much reward as they can before being automatically terminated. We'll return to this concept shortly.

Discount factors are slightly more complicated, and are introduced to model an exponential decay in the value of rewards over time. Concretely, with a discount factor of γ , taking action a_t from state s_t at timestep t and ending up in state s_{t+1} results in a reward of $\gamma^t R(s_t, a_t, s_{t+1})$ instead of just $R(s_t, a_t, s_{t+1})$. Now, instead of maximizing the **additive utility**

$$U([s_0, a_0, s_1, a_1, s_2, \dots]) = R(s_0, a_0, s_1) + R(s_1, a_1, s_2) + R(s_2, a_2, s_3) + \dots$$

we attempt to maximize **discounted utility**

$$U([s_0, a_0, s_1, a_1, s_2, \dots]) = R(s_0, a_0, s_1) + \gamma R(s_1, a_1, s_2) + \gamma^2 R(s_2, a_2, s_3) + \dots$$

Noting that the above definition of a discounted utility function looks similar to a **geometric series** with ratio γ , we can prove that it's guaranteed to be finite-valued as long as the constraint $|\gamma| < 1$ (where $|n|$ denotes the absolute value operator) is met through the following logic:

$$\begin{aligned} U([s_0, s_1, s_2, \dots]) &= R(s_0, a_0, s_1) + \gamma R(s_1, a_1, s_2) + \gamma^2 R(s_2, a_2, s_3) + \dots \\ &= \sum_{t=0}^{\infty} \gamma^t R(s_t, a_t, s_{t+1}) \leq \sum_{t=0}^{\infty} \gamma^t R_{max} = \boxed{\frac{R_{max}}{1 - \gamma}} \end{aligned}$$

where R_{max} is the maximum possible reward attainable at any given timestep in the MDP. Typically, γ is selected strictly from the range $0 < \gamma < 1$ since values in the range $-1 < \gamma \leq 0$ are simply not meaningful in most real-world situations—a negative value for γ means the reward for a state s would flip-flop between positive and negative values at alternating timesteps.

Markovianess

Markov decision processes are "markovian" in the sense that they satisfy the **Markov property**, or **memoryless property**, which states that the future and the past are conditionally independent, given the present. Intuitively, this means that, if we know the present state, knowing the past doesn't give us any more information about the future. To express this mathematically, consider an agent that has visited states s_0, s_1, \dots, s_t after taking actions a_0, a_1, \dots, a_{t-1} in some MDP, and has just taken action a_t . The probability that this agent then arrives at state s_{t+1} given their history of previous states visited and actions taken can be written as follows:

$$P(S_{t+1} = s_{t+1} | S_t = s_t, A_t = a_t, S_{t-1} = s_{t-1}, A_{t-1} = a_{t-1}, \dots, S_0 = s_0)$$

where each S_t denotes the random variable representing our agent's state and A_t denotes the random variable representing the action our agent takes at time t . The Markov property states that the above probability can be simplified as follows:

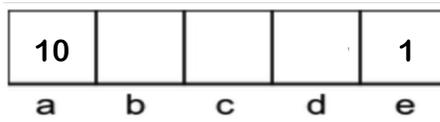
$$P(S_{t+1} = s_{t+1} | S_t = s_t, A_t = a_t, S_{t-1} = s_{t-1}, A_{t-1} = a_{t-1}, \dots, S_0 = s_0) = P(S_{t+1} = s_{t+1} | S_t = s_t, A_t = a_t)$$

which is "memoryless" in the sense that the probability of arriving in a state s' at time $t + 1$ depends only on the state s and action a taken at time t , not on any earlier states or actions. In fact, it is these memoryless probabilities which are encoded by the transition function: $T(s, a, s') = P(s' | s, a)$.

Solving Markov Decision Processes

Recall that in deterministic, non-adversarial search, solving a search problem means finding an optimal plan to arrive at a goal state. Solving a Markov decision process, on the other hand, means finding an optimal **policy** $\pi^* : S \rightarrow A$, a function mapping each state $s \in S$ to an action $a \in A$. An explicit policy π defines a reflex agent - given a state s , an agent at s implementing π will select $a = \pi(s)$ as the appropriate action to make without considering future consequences of its actions. An optimal policy is one that if followed by the implementing agent, will yield the maximum expected total reward or utility.

Consider the following MDP with $S = \{a, b, c, d, e\}$, $A = \{East, West, Exit\}$ (with *Exit* being a valid action only in states a and e and yielding rewards of 10 and 1 respectively), a discount factor $\gamma = 0.1$, and deterministic transitions:



Two potential policies for this MDP are as follows:



With some investigation, it's not hard to determine that Policy 2 is optimal. Following the policy until making action $a = Exit$ yields the following rewards for each start state:

Start State	Reward
a	10
b	1
c	0.1
d	0.1
e	1

We'll now learn how to solve such MDPs (and much more complex ones!) algorithmically using the **Bellman equation** for Markov decision processes.

The Bellman Equation

In order to talk about the Bellman equation for MDPs, we must first introduce two new mathematical quantities:

- The optimal value of a state s , $U^*(s)$ – the optimal value of s is the expected value of the utility an optimally-behaving agent that starts in s will receive, over the rest of the agent's lifetime. Note that frequently in the literature the same quantity is denoted with $V^*(s)$.
- The optimal value of a Q-state (s, a) , $Q^*(s, a)$ - the optimal value of (s, a) is the expected value of the utility an agent receives after starting in s , taking a , and acting optimally henceforth.

Using these two new quantities and the other MDP quantities discussed earlier, the Bellman equation is defined as follows:

$$U^*(s) = \max_a \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma U^*(s')]$$

Before we begin interpreting what this means, let's also define the equation for the optimal value of a Q-state (more commonly known as an optimal **Q-value**):

$$Q^*(s, a) = \sum_{s'} T(s, a, s') [R(s, a, s') + \gamma U^*(s')]$$

Note that this second definition allows us to reexpress the Bellman equation as

$$U^*(s) = \max_a Q^*(s, a)$$

which is a dramatically simpler quantity. The Bellman equation is an example of a *dynamic programming equation*, an equation that decomposes a problem into smaller subproblems via an inherent recursive structure. We can see this inherent recursion in the equation for the Q-value of a state, in the term $[R(s, a, s') + \gamma U^*(s')]$. This term represents the total utility an agent receives by first taking a from s and arriving at s' and then acting optimally henceforth. The immediate reward from the action a taken, $R(s, a, s')$, is added to the optimal discounted sum of rewards attainable from s' , $U^*(s')$, which is discounted by γ to account for the passage of one timestep in taking action a . Though in most cases there exists a vast number of possible sequences of states and actions from s' to some terminal state, all this detail is abstracted away and encapsulated in a single recursive value, $U^*(s')$.

We can now take another step outwards and consider the full equation for Q-value. Knowing $[R(s, a, s') + \gamma U^*(s')]$ represents the utility attained by acting optimally after arriving in state s' from Q-state (s, a) , it becomes evident that the quantity

$$\sum_{s'} T(s, a, s') [R(s, a, s') + \gamma U^*(s')]$$

is simply a weighted sum of utilities, with each utility weighted by its probability of occurrence. This is by definition the *expected utility* of acting optimally from Q-state (s, a) onwards! This completes our analysis and gives us enough insight to interpret the full Bellman equation—the optimal value of a state, $U^*(s)$, is simply the *maximum expected utility* over all possible actions from s . Computing maximum expected utility for a state s is essentially the same as running expectimax—we first compute the expected utility from each Q-state (s, a) (equivalent to computing the value of chance nodes), then compute the maximum over these nodes to compute the maximum expected utility (equivalent to computing the value of a maximizer node).

One final note on the Bellman equation – its usage is as a *condition* for optimality. In other words, if we can somehow determine a value $U(s)$ for every state $s \in S$ such that the Bellman equation holds true for each of these states, we can conclude that these values are the optimal values for their respective states. Indeed, satisfying this condition implies $\forall s \in S, U(s) = U^*(s)$.