

PS3: Practice problems: Revised

1 To Drop Or Not To Drop

A student is trying to decide whether to take Frobology from Professor X or Professor Y, or not to take it at all. There is no homework in either version of the course. They both have a final exam. The only difference between them is that Professor X gives a pass/fail midterm quiz, which is graded before the drop date. In either case, the ultimate outcome is a grade (let's just say, for simplicity, it's A, C, or F). The final grade distributions are the same in both versions. Remember also that it's not necessarily fun to take a quiz.

1. Draw the decision tree for a student faced with this decision. (There is no need to model the decision of whether to drop the course if you take it from Professor Y.)
2. What probabilities and utilities would you need to know in order to compute the best course of action?
3. Describe, in English, a set of circumstances under which Professor X's course would be preferred.
4. Describe, in English, a set of circumstances under which Professor Y's course would be preferred.

2 Short answers

1. Since belief space is real valued (and therefore infinite), why is the branching factor in a POMDP finite?
2. In a POMDP, we can think of the reward of a belief state, b , as being $\sum_s b(s)R(s)$. Why doesn't this lead to self delusion?
3. Why is the finite-horizon optimal policy for an MDP non-stationary, while the infinite-horizon optimal policy is stationary?
4. Give an intuitive motivation for discounting.
5. Consider reinforcement learning. With the same amount of experience, would you expect Q-learning or adaptive dynamic programming (estimate the transition probabilities and rewards, and solve the model) to arrive at a better policy? Under what circumstances (if any) would you prefer the other method?
6. What makes it difficult to do Q learning in very large action spaces?

7. Given that Q can be defined in terms of V and vice versa, why is it important to learn Q rather than V ?
8. Consider a finite-horizon MDP for which we know the initial state. We know two strategies for finding the optimal k -step policy: (1) build a depth- k tree with state s at the root and (2) calculate the optimal k -step policy for all states. Under what circumstances is each approach best?

3 Slightly Less Naive Bayes

In the usual application of naive Bayes to supervised learning, we predict class 1 when $\Pr(Y = 1|X) > \Pr(Y = 0|X)$, where X is the feature vector describing the instance.

What if we were in a situation in which one type of error was much more costly than another? For instance, if we're diagnosing a terrible disease, it might cost much more to miss a diagnosis than to make an erroneous one. So, our utilities might be described as in this table:

Predicted Y	Actual Y	Utility
0	0	0
0	1	-10
1	0	-1
1	1	0

If these are our utilities, then give a rule, in terms of $\Pr(Y = 1|X)$ and $\Pr(Y = 0|X)$, for making predictions that maximize expected utility.

4 Decision Theory

Dr. No has a patient who is very sick. Without further treatment, this patient will die in about 3 months. The only treatment alternative is a risky operation. The patient is expected to live about 1 year if he survives the operation; however, the probability that the patient will not survive the operation is 0.3.

1. Draw a decision tree for this simple decision problem. Show all the probabilities and outcome values.
2. Let $U(x)$ denote the patient's utility function, where x is the number of months to live. Assuming that $U(12) = 1.0$ and $U(0) = 0$, how low can the patient's utility for living 3 months be and still have the operation be preferred?

For the rest of the problem, assume that $U(3) = 0.8$.

3. Dr. No finds out that there is a less risky test procedure that will provide uncertain information that predicts whether or not the patient will survive the operation. When this test is positive, the probability that the patient will survive the operation is increased. The test has the following characteristics:

- True-positive rate: The probability that the results of this test will be positive if the patient will survive the operation is 0.90.
- False-positive rate: The probability that the results of this test will be positive if the patient will not survive the operation is 0.10.

What is the patient's probability of surviving the operation if the test is positive?

4. Assuming the patient has the test done, at no cost, and the result is positive, should Dr. No perform the operation?
5. It turns out that the test may have some fatal complications, i.e., the patient may die during the test. Draw a decision tree showing all the options and consequences of Dr. No's problem.
6. Suppose that the probability of death during the test is 0.005 for the patient. Should Dr. No advise the patient to have the test prior to deciding on the operation?

5 Markov Decision Processes

You are playing a game at a carnival, in which you are trying to throw balls through a hoop. You are allowed to play this game for a total of k steps. If, at the end of k steps, you have gotten at least one ball through the hoop, then you win \$10. If not, you win nothing. On each step, you are allowed to buy, for \$1 each, as many balls as you would like, which you will try to throw simultaneously through the hoop. Each ball has a probability of p of going through the hoop, and each ball's success is independent of the successes of the other balls and of the number of balls being thrown. After you've thrown one set of balls, you can observe whether or not any of them went through the hoop.

1. If you throw n balls at once, what is the probability of getting at least one ball through the net? Write an expression in terms of p and n . Call this quantity $f(p, n)$ in future parts of this problem.
2. If $k = 1$, that is, you can only play one round of this game, what is the optimal number, n , of balls to buy and throw? (You only need to write down an expression involving n and p ; don't worry about getting a closed form).
3. Let s_1 be the state of not having gotten a ball through the hoop and s_2 be the state of having gotten one through. Let $V^k(s)$ be the value of being in state s with k steps remaining in the game. Then $V^0(s_1) = 0$ and $V^0(s_2) = 10$. What is $V^k(s_2)$, assuming there is no discounting.
4. Write an expression for $V^k(s_1)$, in terms of $V^{k-1}(s)$.

6 Tiger hunting: Revised

Let's consider the tiger problem, augmented with two new actions: shoot the tiger on the left and shoot the tiger on the right. Shooting at the door where the tiger is means that, with probability 0.9, the tiger will be scared and move to the other room, and with probability 0.1, he'll stay where he is. Shooting at the door where there is no tiger will have no effect.

Assume that the reward of listening is -1, of shooting is -5, of opening any door without a live tiger is +10, and of opening a door with a live tiger is -50.

There are two possible observations: noise-left and noise-right. If there is a tiger on the right, then $P(\text{noise-right}) = 0.85$, and $P(\text{noise-left}) = 0.15$. The situation is symmetric when the tiger is on the left.

Assume a discount factor of 0.9, but that the game is over when a door is opened.

Draw the alpha vectors for the following short policy trees. Which one is best in which situations?

- Always listen.
- Shoot left, then open left.
- Shoot left, then shoot left, then open left.
- Listen; if noise-left, then open right; if noise-right, then open left.
- Listen; if noise-left, then shoot left, then open left; if noise-right, then shoot right, then open right.
- Listen; if noise-left, then shoot right, then open right; if noise-right, then shoot left, then open left.