

CMPUT 365: Introduction to Reinforcement Learning

Winter 2023

Worksheet #0: Probability Questions

Manuscript version: *#b3d1d6-dirty* - 2025-01-07 22:24:05-07:00

For reviewing probability, either discuss Chapters 1-3 of Evans-Rosenthal ([here](#)), or Martha's ML [notes](#).

Question 1. Suppose that in a lottery you have 0.01 chance of winning and the prize is \$1000. The ticket to enter the lottery costs you \$10. What is the expected amount you would earn, when buying a ticket for this lottery?

Question 2. Adam and Martha propose a simple dice game to you. You can throw a (fair) dice up to two times, and you will take the face value of the last throw in dollars. For example, if you throw the dice once and 3 comes up and you choose to throw again, and, unlucky for you, 2 comes up, you win \$2. Had you decided not to throw again (but of course, you did not know at the time of the decision the outcome of the second throw), you could have taken home \$3.

1. Suppose in your first roll, the dice comes up as a 1. At this stage you stop and ask yourself: How much can I earn on expectation if I roll again?
 2. For what values in your first roll should you re-roll the die – assuming that your goal is to maximize your expected earnings.
 3. What is the expected amount you would earn in this game if you play optimally?
-

Question 3. Let X and Y be discrete random variables and assume that $\mathbb{E}[X]$ exists. Prove the tower property

$$\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X|Y]].$$

Hint: Recall that $\mathbb{E}[X|Y]$ is defined as follows: Let \mathcal{Y} denote the set of values that Y can take with positive probability. By assumption, \mathcal{Y} is at most countably infinite. Define $g : \mathcal{Y} \rightarrow \mathbb{R}$ using $g(y) = \mathbb{E}[X|Y = y]$ ($y \in \mathcal{Y}$). Then, $\mathbb{E}[X|Y]$ is defined using $\mathbb{E}[X|Y] = g(Y)$ (i.e., $\mathbb{E}[X|Y]$ is just the shorthand for $g(Y)$ where g is defined as above).

Now, there is a subtlety in proving this result if both X and Y take on infinitely many values. In this case, you may find yourself wanting to swap two infinite sums. Can we do this? Not, in general! For $i, j \in \mathbb{N} = \{1, 2, \dots\}$, let a_{ij} be defined as follows: $a_{ii} = 1$ for all $i \in \mathbb{N}$. Also, for all $i \in \mathbb{N}$, $a_{i, i+1} = -1$. All the other values are zero. Then, $\sum_i \sum_j a_{ij} = 0 \neq 1 = \sum_j \sum_i a_{ij}$! The numbers $(a_{ij})_{ij}$ need to be “nice” for the two sums to be the same! The result needed here is a special case of the so-called Fubini-Tonelli theorem. For sums, this could be phrased as follows: Consider an infinite array $(a_{ij})_{i, j \in \mathbb{N}}$. If either $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}$, or $\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij}$ exist then both exist and they are equal.

Question 4. Consider the following game: you roll two (fair) 6-sided dice and win \$1 if the sum of the dice roll is 2, 5, 7, 8 or 11. Otherwise, you lose \$1.

1. What is the expected value of the sum of the outcomes of the two dice rolls?
 2. What is the variance of the sum of the outcomes of the two dice rolls?
 3. What is the expected value of the winnings for playing this game? In other words, how much money are you expected to gain (or lose).
 4. What is the variance of the winnings for this game?
 5. Would you play this game as stated above? Only a single shot. How about if the amount won or lost was \$100? How about \$1000?
 6. What if I say let's play this game a 1000 times. Of course, every play is independent of the previous ones. Does this change what you say? In which case?
 7. Let's say we also add 12 to the list of numbers when you win. How does this change your answer to the last two questions if at all?
-

Question 5. What is the maximum likelihood estimate (MLE) of parameter $\theta \in D := (-1, \infty)$ given the observed data $D = (y_1, y_2, \dots, y_{n-1}, y_n) \in \mathbb{R}^n$ and assuming the density model $f_\theta(y) = (\theta + 1)y^\theta, 0 < y < 1$ for the data, and independent observations?

Question 6. Now suppose a game where you choose to flip one of two (possibly unfair) coins. You win \$1 if your chosen coin shows heads and lose \$1 if it shows tails. Note that you do not know the probability of the coin outcomes. Instead, you are able to view 6 sample flips for each coin respectively: (T, H, H, T, T, T) and (H, T, H, H, H, T) .

1. For each coin: what is the MLE estimate for the probability of heads p_1 (respectively, p_2)?
2. Using Hoeffding's inequality (see below), construct two confidence intervals $[a_1, b_1]$ and $[a_2, b_2]$ such that with probability 0.8 or higher, both $p_1 \in [a_1, b_1]$ and $p_2 \in [a_2, b_2]$ hold. The intervals should not be wider than what the tightest interval based on Hoeffding's inequality can be. Calculate the end-points up to 3 digits and use the usual rules of rounding when showing the results.
3. Which coin would you flip? Would you be willing to flip a coin other than the one you chose?

Hint: Hoeffding's inequality has many forms, but the following one will be particularly well suited for our purposes: Let $X_1, \dots, X_n \in [a, b]$ be independent random variables with common mean (=expected value) $\mu \in [a, b]$. Let $\hat{\mu} = (X_1 + \dots + X_n)/n$. Then, for any $\delta \in [0, 1]$, with probability $1 - \delta$,

$$\mu \leq \hat{\mu} + (b - a) \sqrt{\frac{\log(1/\delta)}{2n}}.$$

Question 7. We say that an estimator (a function mapping data to reals) is **unbiased** if its expected value on the data equals the “value to be estimated”. Assume, for example that our goal is to estimate the value of an integral of a function $f : \mathbb{R} \rightarrow \mathbb{R}$ with respect to a complicated density function $p : \mathbb{R} \rightarrow [0, \infty)$ and at our disposal is a function that takes a real $x \in \mathbb{R}$ and returns the values $p(x), f(x)$ (e.g., in your favorite computer language).

$$u = \int f(x)p(x)dx = ?$$

We assume that μ is well-defined.

One possibility is to use a numerical technique (quadrature rule, etc.) to approximately calculate the value of μ . Another possibility is to use a “Monte Carlo” approach: Let $q : \mathbb{R} \rightarrow [0, \infty)$ be another density function and assume that if $p(x) > 0$ for some $x \in \mathbb{R}$, then $q(x) > 0$ also holds (in cases like this, with a figure of speech, we say that p is “dominated by” q). Further assume that we have another function at our disposal that implements sampling from q and it can also return the value that q takes at any point.

1. Show that $R = \frac{1}{n} \sum_{i=1}^n \frac{p(X_i)}{q(X_i)} f(X_i)$ gives an unbiased estimate of μ , where X_1, \dots, X_n are independent draws from q . (Note that R can be easily calculated given the two functions at our disposal by issuing n calls to these functions.)
2. Show that $\mathbb{V}[R] = \frac{1}{n} \sigma(q)$ where $\sigma(q) = \int_{q>0} \frac{f^2(x)p^2(x) - \mu^2 q^2(x)}{q(x)} dx$, which may or may not be finite.
3. (**)¹ Sometimes we have a choice in what q to use. When we do, we usually aim for getting a small value for $\sigma(q)$. Sometimes this does not work though: Give an example when $\sigma(q) = \infty$. How will the estimate R behave in this case? Can we trust it?
4. Assume that f is nonnegative valued. Show that the variance of R is minimized by choosing q as $q^*(x) = f(x)p(x)/\mu$.
5. Is it realistic to assume that we will have a function to sample from q^* and to evaluate q^* ?
6. (***) What density minimizes the variance of R in the general case when f can take on any real value, including negative ones?

¹This means, this is an advanced level question. The more stars, the harder the question is.