Algorithms M2–IF TD 1

October 6, 2021

1 Big-O notation

Consider the following functions:

$$2^n, 3^n, n^3, n \log n, n!, n^n, 2^{\sqrt{\log n}}, 3^{\log n}, 2^{\log \log n}, \log^2 n, (\log n)^{\log n}, \log(n!), 2^{\sqrt{n}}, 2^{n^2}$$

Sort them in increasing order of Big-O notation.

Solution:
In increasing order we have the functions

$$2^{\log \log n}, \log^2 n, 2^{\sqrt{\log n}}, n \log n \approx \log(n!), 3^{\log n}, n^3, (\log n)^{\log n}, 2^{\sqrt{n}}, 2^n, 3^n, n!, n^n, 2^{n^2}$$

2 Basic Probabilities

[MU Ex. 1.1] We flip a fair coin ten times. Calculate the following probabilities:

1. The number of heads is equal to the number of tails.
2. There were more heads than tails. (What if we flip 11 times?)
3. The 2nd and 9th flip produced the same outcome.
4. We flipped at least eight consecutive heads.

Solution:

1. We have $p = \binom{10}{5} 2^{-10}$. Recall that $\binom{10}{5} = \frac{10!}{5!5!}$ and generally $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ denotes the number of ways we can select k elements out a set of size n.
2. Let p_1 be the probability that we have more heads, p_2 be the probability that we have more tails, and p_3 be the probability of a tie (which we calculated in the previous question). Then $p_1 = p_2$ by symmetry. Also, $p_1 + p_2 + p_3 = 1$. So, $p_1 = \frac{1-p_3}{2}$.
3. $\frac{1}{2}$.

4. Let A_1 be the event that flips 1 to 8 are all heads, A_2 the event that flips 2 to 9 are all heads, A_3 the event that flips 3 to 10 are all heads. The probability we are interested in is $p = Pr[A_1 \cup A_2 \cup A_3] = Pr[A_1] + Pr[A_2] + Pr[A_3] - Pr[A_1 \cap A_2] - Pr[A_2 \cap A_3] - Pr[A_1 \cap A_3] + Pr[A_1 \cap A_2 \cap A_3]$ (this is called the inclusion-exclusion principle). We now have $Pr[A_i] = 2^{-8}$, $Pr[A_1 \cap A_2] = Pr[A_2 \cap A_3] = 2^{-9}$ and $Pr[A_1 \cap A_3] = Pr[A_1 \cap A_2 \cap A_3] = 2^{-10}$. So in the end the probability of 8 consecutive heads is $3 \cdot 2^{-8} + 2 \cdot 2^{-9} = 2^{-7}$.

3 Basic Probabilities 2: Modulo arithmetic

[MU Ex. 1.15] We roll ten standard dice (six-sided). What is the probability that the sum of our dice is divisible by 6?

Solution:
$\frac{1}{6}$. To see this, suppose that the sum of the first nine dice is x and that the last die rolled a $y \in \{1, 2, 3, 4, 5, 6\}$. For any value of x, there exists exactly one value of y that makes $x + y$ a multiple of 6.

4 Basic Probabilities 3: Independence

Give an example of three events A, B, C which are pair-wise independent but not independent as a set.

Solution:
We flip two coins. A is the event that the first coin is heads, B is the event that the second coin is heads. C is the event that the two coins are equal. We have $Pr[A] = Pr[B] = Pr[C] = \frac{1}{2}$, $Pr[A \cap B] = Pr[A \cap C] = Pr[B \cap C] = \frac{1}{4}$. However, $Pr[A \cap B \cap C] = \frac{1}{4}$. This proves that the three events are not all independent, since if they were, we would have $Pr[A \cap B \cap C] = Pr[A] \cdot Pr[B] \cdot Pr[C]$.

5 Generating Randomness

Suppose that we are given a fair coin. Explain how to use repeated flips of this coin to:

1. Generate a number in $\{2, \dots, 12\}$ which has the same distribution as the sum of two normal dice.

2. Generate a random number in $\{0, \dots, 10\}$.

How many flips will you need (in expectation)?

Suppose that we are given a non-fair coin that has probability $p \neq \frac{1}{2}$ of giving heads. Show how you can use this coin to generate a random bit with probability $\frac{1}{2}$ of being 1. How many coin flips do you need (in expectation)?

Solution:

Let us first explain how to generate a number uniformly at random in $\{0, \ldots, x\}$ for any positive integer x. Let q be the smallest integer such that $2^q > x$. Flip the coin q times. There are $2^q \geq x + 1$ possible outcomes. For each value $i \in \{0, \ldots, x\}$, we assign it a distinct outcome (out of the 2^q possible outcomes), and we output i if this outcome happened. If the result of the coin flips does not correspond to a value, we repeat. The probability that we output something is at least $\frac{1}{2}$, because $x \geq 2^{q-1}$, so the expected number of repetitions of this experiment is 2, given an expected number of $2q \approx 2 \log x$ coin flips.

This covers the generation of a number in $\{0, \ldots, 10\}$. For simulating the dice, we do the same by producing two numbers in $\{1, \ldots, 6\}$, and taking their sum.

For the second part, flip the coin twice. If the outcome is HT output 0. If it is TH output 1. Otherwise, restart the experiment. The probability of success is $2p(1 - p)$, so the expected number of repetitions is $\frac{1}{2p(1-p)}$.

6 False-Positive Paradox

Algorithmitis is a rare disease that affects roughly $\frac{1}{100}$ of the general population. Prof. Chaos has come up with an innovative new test to detect this disease in the population. This test has a success rate of $\frac{98}{100}$, but possibly two-sided error: with probability $\frac{2}{100}$ it may misidentify a sick person as healthy, or a healthy person as suffering from algorithmitis.

We randomly select a volunteer from the general population and perform the test on her. Suppose that the test came back positive. What is the probability that the volunteer has algorithmitis?

Solution:

You may naively believe that the answer is $\frac{98}{100}$, since this is the success rate of the test. IT IS NOT!

Let S be the event that the patient tested is sick. Let D be the event that the test detected the sickness. We know that $Pr[D \mid S] = \frac{98}{100}$ and that $Pr[\neg D \mid \neg S] = \frac{98}{100}$, that is, if the patient is sick, with probability 0.98 we detect it, and if the person is not sick, we detect that she is healthy with probability 0.98.

However, this exercise asks for the probability $Pr[S \mid D]$: what is the probability that someone is sick, if the test said so. This is NOT the same as $Pr[D \mid S]$. We have $Pr[S \cap D] = Pr[S \mid D] \cdot Pr[D] = Pr[D \mid S] \cdot Pr[S]$. We know $Pr[D \mid S] = \frac{98}{100}$ and $Pr[S] = \frac{1}{100}$. So, all we are missing is $Pr[D]$, which is the probability of a random member of the population being detected as sick by the test (whether she is sick or not).

We have $Pr[D] = Pr[D \mid S] \cdot Pr[S] + Pr[D \mid \neg S] \cdot Pr[\neg S] = \frac{98}{100} \cdot \frac{1}{100} + \frac{2}{100} \cdot \frac{99}{100}$. Plugging this into the equation above gives $Pr[S \mid D] = \frac{98}{296} \approx \frac{1}{3}$. So, even though the test seems super-accurate, even when it is positive, it's still more likely than not that the patient is healthy!

7 Monty Hall Paradox

In the game show "Let's make a deal" the contestant is given three closed boxes A,B,C. One of the boxes contains a prize, while the others are empty. The game is played as follows:

1. The contestant picks a box, say A.
2. The host considers the remaining boxes B and C. He publicly opens one of the two, always selecting to open a box that is empty. **NB:** This is always possible (why?)
3. The contestant must either hold his box or exchange with the remaining un-opened one. No other information is given.

What is the probability that the contestant wins if she decides to exchange her box?

Solution:
$\frac{2}{3}$. This is called a paradox because many people would think that the contestant did not acquire any new information when the host opened a box (we already knew that one of the remaining boxes was empty). However, the box the host opened was not random! In the beginning, there was a $\frac{2}{3}$ probability that B or C have the prize. This probability did not change, because the host made sure to keep the box with the prize closed, if it was one of B, C.

8 Min Cut

For the Min Cut algorithm that we saw in class, suppose that we make the following modification: at each step, instead of picking a random edge (u, v) to contract, the algorithm picks two random vertices $u, v \in V$ and collapses them into a single vertex. Prove (by giving a counter-example) that this algorithm has, in some instances, exponentially small probability of finding a minimum cut.

Solution:
Take two cliques of size $n/2$ and add an edge connecting them. The min cut is clearly this new edge. However, with probability $1/2$ we will collapse two vertices from different cliques. If we do this, the min cut size increases. The probability that this does not happen in k consecutive steps is at most "roughly" (but not exactly) 2^{-k}. This counter-example shows why it is crucial for the Min Cut algorithm to only contract nodes which are connected by an edge: if we contract a random pair, the probability that the two nodes "should" be on different sides of the cut (and therefore we made a mistake) is $1/2$, so the probability of success drops exponentially with every round. However, if we select the endpoints of an edge, we can take into account that most edges are **not** cut in the optimal solution, so "probably" (with probability $1 - \frac{1}{poly(n)}$) the nodes we merge should be on the same side of the cut.