15-251: Great Theoretical Ideas In Computer Science

Recitation 9 Solutions

Conditions permitting

(a) Mary flips a fair coin. If it lands heads, she rolls a 3-sided die; otherwise she rolls a 4-sided die. What is the probability that she rolls at least a 3?

(b) Martha is making three pancakes, each of which burns with probability 0.5. Given that at least one pancake burns, what is the probability that they all burn?

(a) We can draw the "outcome tree" and note that the three outcomes in the event we care about are (H,3), (T,3), and (T,4), which occur with probabilities $\frac{1}{6}$, $\frac{1}{8}$, and $\frac{1}{8}$ respectively. So the probability of this event is their sum.
(b) We can draw the probability tree again to see that there are eight equiprobable outcomes, of which seven involve at least one burnt pancake. One of these seven outcomes is the one where all pancakes burn, so the probability we're looking for is $\frac{1}{7}$.

T totals

(a) The law of total probability states that any three events A, B, T where A, B partition the sample space ($A \cap B = \emptyset$ and $A \cup B = S$),

$$Pr[T] = Pr[T|A] \cdot Pr[A] + Pr[T|B] \cdot Pr[B]$$

Prove this fact and note that it extends to partitions of any size.

(b) Alice flips a fair coin $n + 1$ times and Bob flips a fair coin n times. What is the probability that Alice's coin comes up heads strictly more than Bob's does?

(a)

$$Pr[T|A] \cdot Pr[A] + Pr[T|B] \cdot Pr[B] = \frac{Pr[T \cap A]}{Pr[A]} \cdot Pr[A] + \frac{Pr[T \cap B]}{Pr[B]} \cdot Pr[B]$$
$$= Pr[T \cap A] + Pr[T \cap B]$$
$$= \sum_{x \in T \cap A} p(x) + \sum_{x \in T \cap B} p(x) = \sum_{x \in T} p(x) = Pr[T]$$

(b) Let A be the number of heads Alice gets and A' be the number of heads she gets on the first n flips. Let B be the number of heads Bob gets and let L be the event that Alice's last flips lands heads. Using total probability with the partition $(L, \overline{L})$, we can see

$$Pr[A > B] = Pr[A > B \mid L] \cdot Pr[L] + Pr[A > B \mid \overline{L}] \cdot Pr[\overline{L}]$$
$$= Pr[A' \geq B] \cdot \frac{1}{2} + Pr[A' > B] \cdot \frac{1}{2}$$

By symmetry, note that $Pr[A' > B] = Pr[A' < B]$. So the above is equal to

$$Pr[A' \geq B] \cdot \frac{1}{2} + Pr[A' < B] \cdot \frac{1}{2}$$
$$= \frac{1}{2}(Pr[A' \geq B] + Pr[A' < B]) = \frac{1}{2}$$

Dog gone

You lose sight of your dog running in the park one day. When you chase after, you come to a fork. Each branch of the fork goes on infinitely far, but your dog chose one of the branches, walked d meters down it, and stopped. Describe a good algorithm to find your dog and analyze its competitive ratio. (Cost is the total distance you walk.)

Note that the optimal offline algorithm knows which side the dog went to and just walks d meters down the correct path.
Online algorithm: walk one meter down the right side, then come back. Walk one meter down the left side, then come back. Repeat, doubling your search distance each time.
The worst case for this algorithm is where the dog walked $2^k + 1$ meters down the left path for some $k \in \mathbb{N}$. The distance you walk in this case is $(\sum_{i=0}^{k} 4 \cdot 2^i) + 2 \cdot 2^{k+1} + 2^k + 1$, which comes out to approximately $4 \cdot 2^{k+1} + 2 \cdot 2^{k+1} + 2^k$ which is only about $13d$ (recall $d = 2^k + 1$).
Can you think of a way to improve the competitive ratio?
Fun fact: you can improve this strategy a little more using randomization. (How?)

Blunt truths

Your friend Norville is showing off his fair 7-sided die and wants to play a game where you keep rolling the die until you roll a 7. He claims that the game is more likely to end after an odd number of rolls "because 7 is, like, odd, man." Is Norville right?

The probability of the game ending after exactly r rolls is $\frac{1}{7} \cdot (\frac{6}{7})^{r-1}$ (we can formally prove this using the chain rule). Let A be the event that the game ends after an odd number of rolls and use total probability to see that $Pr[A] = \frac{1}{7}\sum_{i=0}^{\infty}(\frac{6}{7})^{2i}$. Similarly, let B be the event that the game ends after an even number of rolls and note that $Pr[B] = \frac{1}{7}\sum_{i=0}^{\infty}(\frac{6}{7})^{2i+1}$. Now subtract the two and see that

$$Pr[A] - Pr[B] = \frac{1}{7}\sum_{i=0}^{\infty}(\frac{6}{7})^{2i} - (\frac{6}{7})^{2i+1}$$

Each term in the sum is positive, so the whole thing is strictly positive, meaning that A is more likely than B after all.

Online dating

There are n people with whom you can go on dates, and you would like to propose marriage to the best one. Of course, once you've proposed to somebody, it would be rude to keep going on dates with other people, but you can only propose to the last person you went on a date with. So after each date, you must immediately decide whether or not to propose to that person. (We assume that you can compare people and that no two people are exactly equally good.)

Your strategy is as follows: first, you go on dates with t people and reject them. Then, continue dating until you meet somebody who is better than all of the first t, and propose to that person. (If nobody meets this criterion, remain forever alone.)

For what value of t do you maximize the probability that you propose to the best candidate?

You propose to the best person if and only if the first date with value greater than the maximum of the first t candidates is the best. Thus we can express the probability of selecting the best person as follows:

$\sum_{j=t+1}^{n} P(\text{we pick the best candidate} \mid \text{candidate } j \text{ is best})P(\text{candidate } j \text{ is best})$.

Then, the probability that candidate j is the best is simply $\frac{1}{n}$, and the probability that we pick candidate j given that it is best is exactly the probability that among the first $j-1$ candidates, the maximum was in the first t candidates. The maximum is equally likely to be in any of the $j-1$ spots, so this has probability $\frac{t}{j-1}$. Thus the probability of picking the best person is

$\frac{t}{n}\sum_{j=t+1}^{n}\frac{1}{j-1} = \frac{t}{n}(\sum_{j=1}^{n-1}\frac{1}{j} - \sum_{j=1}^{t-1}\frac{1}{j}) \approx \frac{t}{n}(\text{ln}n - \text{ln}t) = \frac{t}{n}\text{ln}\frac{n}{t}$.

This is maximized at $t = \frac{n}{e}$.