Randomized Algorithms **NYU, Fall 2025**
HW #0 (out Monday Sept 8, 2025) **Due:** Monday Sept 15, 2025

This HW is going out at the end of the first week of classes so that you can get a feel for the course asap, and prepare accordingly. It's a HW with a short deadline (in a week). There are some exercises (which are only for practice), and some actual problems. These problems are solvable using ideas from basic probability and algorithms courses.

Please solve the problems without collaboration. Submissions will be via gradescope, and the link will appear on the course webpage and on Brightspace. Also, changes, corrections, and clarifications will also appear on ED, so please check it regularly.

Exercises

Exercises are for fun and edification, please do not submit. (You may discuss these exercises with others.) The ones below are grouped by topic and their subparts do not necessarily build on one another (for example, you do not need to do (a) to do (b)).

1. (a) You flip n independent unbiased coins—i.e., each coin is 1 (or heads, or true) with probability $1/2$, and 0 (or tails, or false) w.p. $1/2$. We'll be very explicit in this exercise about the sample space, events, etc.
   i. What is the sample space Ω in this problem. How many outcomes are there? What is the probability of each event $\omega \in \Omega$?
   ii. Suppose we define random variable X to be $X(\omega) =$ the number of 1's in ω. What is $\mathbb{E}[X]$?
   iii. Suppose we consider the events $\mathcal{E}_i := \{\omega \mid \omega \text{ contains } i \text{ 1's}\}$ (or equivalently, that $\mathcal{E}_i = \{\omega \mid X(\omega) = i\}$ or more compactly, $\mathcal{E}_i = \{X = i\}$) then what is the probability of event $\mathcal{E}_i$? Argue that these events are disjoint.
   iv. How do these answers change if the bias of the coins becomes $p \in [0, 1]$?

   (b) Given random variables (r.v.s) X, Y, use the definition of the mean and variance to show the following:
   i. $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$ and $\mathbb{E}[cX] = c\mathbb{E}[X]$ and $\mathbf{Var}(cX) = c^2\,\mathbf{Var}(X)$ for any constant c.
   ii. If X, Y are independent, then show that $\mathbb{E}[XY] = \mathbb{E}[X]\mathbb{E}[Y]$ and $\mathbf{Var}(X + Y) = \mathbf{Var}(X) + \mathbf{Var}(Y)$. (Why do you need independence? Can you give non-independent r.v.s X, Y for which these statements are false?)
   iii. For independent $X_1, \ldots, X_n$, if each X_i has mean μ and variance σ^2, then show $\frac{\sum_{i=1}^n X_i}{n}$ has mean μ and variance σ^2/n.

   (c) We have an algorithm A such that each time it is run, it crashes (independently of the past) with probability p. What is the expected number of times it is run until it does not crash?

   (d) Suppose X_1, X_2 are two independent and uniformly random bits, i.e., each of them is 0 with probability (w.p.) 1/2 and 1 w.p. 1/2. Define $X_3 = X_1 \oplus X_2$ to be their XOR.
   i. Show that any pair X_i, X_j of these three random variables are pairwise independent.
   ii. Show that X_1, X_2, X_3 are not collectively independent.

(e) Let X be a uniformly random number in $\{1, 2, \ldots, n\}$. What is $\mathbb{E}[X]$? $\mathbb{E}[X^2]$? Is $(\mathbb{E}[X])^2$ equal to $\mathbb{E}[X^2]$? Similarly, check that $\mathbb{E}[1/X]$ is not equal to $1/\mathbb{E}[X]$. If f is a convex function, what can you say about $f(\mathbb{E}[X])$ vs $\mathbb{E}[f(X)]$? (Remember Jensen's inequality?)

(f) Suppose we pick two independent and uniformly random n-bit strings X and Y, and let $Z = \langle X, Y \rangle \pmod 2$ be their inner product modulo 2. (We use $\langle x, y \rangle := \sum_i x_i y_i$ in this course to denote the inner product.) ~~Show that $\mathbf{Pr}[Z = 1] = 1/2$.~~ Show that $\mathbf{Pr}[Z = 1] = 1/2(1 - 2^{-n})$.

2. (a) Given a tree $T = (V, E)$ and a set of nodes $X \subseteq V$ that contains all the leaves, prove that the average degree of nodes in X is less than 2.

   (b) Suppose graph G has integer weights in the range $\{1, \ldots, W\}$, where $W \geq 2$. Let G_i be the edges of weight at most i, and κ_i be the number of components in G_i. Then show that the MST in G has weight exactly $n - W + \sum_{i=1}^{W-1} \kappa_i$.

   (c) You are give a bipartite graph $G = (U, V, E)$ with $|U| = |V| = n$ and maximum degree Δ. Give ~~a linear-time~~ an $O(n\text{poly}(\Delta))$-time algorithm to color the edges of G with 2Δ colors so that the edges incident to every vertex have distinct colors.

   (d) You want to partition a given input string into segments in the cheapest way. Let $A = a_1, a_2, \ldots, a_n$ be a string. Let $C_{i,j}$ (for $i \leq j$) be the "cost" of a segment $a_i, \ldots a_j$. Assume that these costs $C_{i,j}$ have been precomputed and given to you as inputs, so you can look them any particular cost in unit time. The total cost of some segmentation is the sum of the costs of all segments. (E.g., if you split A into three parts $[a_1, \ldots, a_p][a_{p+1}, \ldots, a_q][a_{q+1}, \ldots, a_n]$ then your cost would be $C_{1,p} + C_{p+1,q} + C_{q+1,n}$. Of course, you can split into as many parts as you like.) Use dynamic programming to compute the least-cost segmentation of A in time $O(n^2)$.

Problems

Please write short and clear solutions to each of these problems. Use the language of probability to your advantage. Be clear what the events are, what probabilities and expectations you are reasoning about.

Probability Problems.

1. You want to sample uniformly at random from the set of all n-bit strings that are *balanced*, i.e., that contain exactly $n/2$-many 0's and $n/2$-many 1's (assume n is even). You do the following: Sample a uniform random string ω from the set of all n-bit strings, and output ω (and stop) if ω is balanced, else reject and sample again.

   Show that the expected number of n-bit strings you sample is $O(\sqrt{n})$. (Hint: You may use Stirling's approximation.)

2. Let $x_1, x_2, \ldots, x_n$ be a random permutation of the numbers $[n] := \{1, 2, \ldots, n\}$. You scan the numbers from left to right. At any time, you maintain a number M in your hand, initially $M = -\infty$. When you see x_i, if $x_i > M$ then set $M \leftarrow x_i$ ("you change M"), else leave M unchanged. Show that the expected number of times M is changed is $H_n := 1 + 1/2 + 1/3 + \ldots + 1/n$.

3. Two judges are both giving opinions on a series of n questions. On each question, they answer True or False (T or F for short). On the first question they each independently answer T or

F with probability $1/2$. Thereafter, they steadily become more "stick-in-the-mud": for the i^{th} question, they independently repeat their own previous answer with probability $1 - \frac{1}{i+1}$, and give the opposite answer with probability $\frac{1}{i+1}$. What is the expected number of answers on which they agree?

Algorithms Problems.

1. Let $A[1 \ldots n]$ be an array of pairwise distinct numbers. A pair (i, j) of indices $1 \leq i < j \leq n$ is called an *inversion* of A if $A[i] > A[j]$. (For example, the array $\langle 8, 5, 2, 7, 9 \rangle$ has four inversions—$(1, 2), (1, 3), (1, 4)$, and $(2, 3)$.)

   Give a deterministic algorithm running in time $\Theta(n \log n)$ for computing the number of inversions in A. You may assume that n is a power of two. (E.g., you could use divide-and-conquer.)

2. You are given a (connected) directed acyclic graph $G = (V, E)$ where each edge $e \in E$ has length ℓ_e. Give a polynomial-time algorithm to find the longest directed path in the graph. Ideally your algorithm should run in linear time in the number of edges of G.