

Assignment #6

Due: 6:00pm Wednesday 2 November 2022

Upload at: <https://www.gradescope.com/courses/438130/assignments/2352912>

Assignments in COS 302 should be done individually. See the [course syllabus](#) for the collaboration policy.

Remember to append your Colab PDF as explained in the first homework, with all outputs visible.
When you print to PDF it may be helpful to scale at 95% or so to get everything on the page.

Problem 1 (15pts)

Consider a sample space $\Omega = \{\text{red, green, blue, orange, yellow}\}$.

- (A) What is the smallest possible valid event space \mathcal{A} ?
- (B) What is the smallest possible event space that contains the set $\{\text{blue}\}$?
- (C) What is the smallest possible event space that contains both $\{\text{blue}\}$ and $\{\text{red, green}\}$?
(Hint: it has eight members.)

Problem 2 (15pts)

Consider the following bivariate distribution $p(x, y)$ of two discrete random variables X and Y .

	y_1	0.01	0.02	0.03	0.1	0.1
	y_2	0.05	0.1	0.05	0.07	0.2
Y	y_3	0.1	0.05	0.03	0.05	0.04
		x_1	x_2	x_3	x_4	x_5
				X		

- (A) What is the marginal distribution $p(x)$?
- (B) What is the marginal distribution $p(y)$?
- (C) What is the conditional distribution $p(x | Y = y_1)$?
- (D) What is the conditional distribution $p(y | X = x_3)$?
- (E) What is the conditional distribution $p(x | Y \neq y_1)$?

Problem 3 (20pts)

In 2014-2016, West Africa experienced a massive outbreak of Ebola. We'll concentrate on Sierra Leone and imagine a mandatory screening of every citizen. The probability of being tested positive given that the citizen has Ebola is 84%. The probability of being tested positive given that the citizen does not have Ebola is 11%. We also know that the probability of contracting Ebola for any given citizen is 0.4%. If a randomly-chosen citizen tests positive, what is that citizen's probability of actually having Ebola?

Problem 4 (48pts)

In this problem you will do some mathematical calculations and also use Colab. Be sure to append your PDF and insert your link as usual.

- (A) Import `numpy.random` (usually aliased to `np.r`) and **set the random seed**.
- (B) Imagine drawing 1,000 independent **Bernoulli variates** $X_i \in \{0, 1\}$ with the probability $p(X_i = 1) = 0.35$, and computing their sum

$$Y = \sum_{i=1}^{1000} X_i .$$

What are the mean and variance of Y ?

- (C) Generate 10,000 independent random variables Y as described above. That is, generate 10,000 sums of 1,000 independent Bernoulli variables. This is not as hard as it sounds. Use the `numpy.random.rand()` function only; do not use any functions from `scipy.stats`. Use `rand()` to generate a $10,000 \times 1,000$ matrix of independent uniform random variates in the interval $[0, 1]$, then threshold them appropriately to get 0 or 1. Finally, sum over the appropriate dimension to get 10,000 samples of Y above.
- (D) Use Matplotlib to **make a histogram** of these samples. Use at least 100 bins so you can see the structure in the distribution. Do you recognize the shape?
- (E) Compute the empirical mean and variance of the 10,000 samples you have drawn. Compare these results to your calculation from (B).
- (F) Now imagine drawing 1,000 independent (continuous) variates uniformly in the interval $[-1, 1]$. What are the mean and variance of their sum?
- (G) Generate 10,000 such sums using a variation of the procedure you performed for (C). As before, you'll only use `rand()` but you should scale and shift rather than threshold.
- (H) Create a histogram of these sums, as in (D). Describe the shape.
- (I) Compute the empirical mean and variance of the 10,000 samples and compare them to your computations in (F).

Problem 5 (2pts)

Approximately how many hours did this assignment take you to complete?

My notebook URL: <https://colab.research.google.com/XXXXXXXXXXXXXXXXXXXXXXXXX>

Changelog

- 12 Oct 2022 – F22 version.