

Assignment #9

Due: 6:00pm April 16, 2021

Upload at: <https://www.gradescope.com/courses/214271/assignments/1156197>

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)

A coin is weighted so that its probability of landing on heads is 0.2, suppose the coin is flipped 20 times.

- (A) Compute the probability that the coin lands on heads at least 16 times.
- (B) Use Markov's inequality to bound the event that the coin lands on heads at least 16 times.
- (C) Use Chebyshev's inequality to bound the event that the coin lands on heads at least 16 times.

Problem 2 (28pts)

The moment generating function (MGF) for a random variable X is:

$$M_X(t) = \mathbb{E}[e^{tX}].$$

One useful property of moment generating functions is that they make it relatively easy to compute weighted sums of independent random variables:

$$Z = \alpha X + \beta Y$$

$$M_Z(t) = M_X(\alpha t)M_Y(\beta t).$$

- (A) Derive the MGF for a Poisson random variable X with parameter λ . (Hint: Note that $e^{tX} = (e^t)^X$ and apply the “pattern matching normalization constant” trick from the previous homework.)
- (B) Let X be a Poisson random variable with parameter λ , as above, and let Y be a Poisson random variable with parameter γ . X and Y are independent. Use the MGF to show that their sum is also Poisson. What is the parameter of the resulting distribution?
- (C) Derive the MGF for an exponential random variable X with parameter λ . For what values of t does the MGF $M_X(t)$ exist?
- (D) Let X be an exponential random variable with parameter λ , and let $\alpha > 0$ be a constant. Show that the random variable αX is also exponentially distributed. What is its parameter?

Problem 3 (30pts)

In this problem, you will use the entire 50k-digit MNIST data set. To remind you, the data are 28×28 greyscale images of the digits 0 through 9. Download the `mnist_full.pkl.gz` file and load it with something like this:

```
import pickle as pkl
import numpy as np
import gzip
with gzip.open('mnist_full.pkl.gz', 'rb') as fh:
    mnist = pkl.load(fh)
```

This will give you a `dictionary` `mnist` with keys and values that should be self-explanatory. It is possible that when you run the code to solve the problems below, that you run into issues in which the covariance matrix is not positive definite. To solve this, add a little bit of diagonal to it, i.e., add $\alpha \mathbb{I}$, for $\alpha \approx 10^{-6}$.

- (A) Compute the empirical mean μ and covariance Σ of the training images.
- (B) Reshape and render the mean as an image using `imshow`.
- (C) Generate 5 samples from the multivariate Gaussian with parameters μ and Σ from (A). Do this only using `numpy.random.randn` and linear algebra operations. Reshape and render these samples using `imshow`.
- (D) Now iterate over each of the possible labels from 0 to 9. Compute the mean and covariance of the training data that have that label. Here's a sketch of some code for getting the correct images:

```
for label in range(10):
    indices = train_labels == label
    images = train_images[indices, :]
```

Render the mean of each as an image, and generate 5 samples from the Gaussian distribution with the label-specific mean and covariance. Render these samples.

Problem 4 (25pts)

Consider the following joint distribution over 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		

Use Colab to compute the following quantities. Remember to compute these in **bits** which means using base 2 for logarithms. Here's the PMF in Python to make life a bit easier:

```
PXY = np.array([[0.01, 0.02, 0.03, 0.1, 0.1],  
                [0.05, 0.1, 0.05, 0.07, 0.2],  
                [0.1, 0.05, 0.03, 0.05, 0.04]])
```

- (A) What is the **entropy** $H(X)$?
- (B) What is the **entropy** $H(Y)$?
- (C) What is the **conditional entropy** $H(X | Y)$?
- (D) What is the **conditional entropy** $H(Y | X)$?
- (E) What is the **joint entropy** $H(X, Y)$?
- (F) What is the **mutual information** $I(X; Y)$? Compute it once using the PMF and then once using relationships between the quantities you used in (A)-(E) to verify your answer.

Problem 5 (2pts)

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

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

Changelog

- 6 April 2021 – S21 Initial version