CIS 4190/5190: Applied Machine Learning — Spring 2025

Homework 3

Handed Out: March 19 — *Due: 8 pm April 2*

- You are encouraged to format your solutions using LaTeX. Handwritten solutions are permitted, but remember that you bear the risk that we may not be able to read your work and grade it properly — we will not accept post hoc explanations for illegible work. You will submit your solution manuscript for written HW 3 as a single PDF file.
- The homework is **due at 8 PM** on the due date. We will be using Gradescope for collecting the homework assignments. Please submit your solution manuscript as a PDF file via Gradescope. Post on Ed Discussion and contact the TAs if you are having technical difficulties in submitting the assignment.
- Make sure to assign pages to each question when submitting homework to Gradescope. The TA may deduct 0.2 points per sub-question if a page is not assigned to a question.

1 Written Questions

Note: You do not need to show work for multiple choice questions. If formatting your answer in LaTeX, use our LaTeX template `hwtemplate`. (This is a read-only link. You'll need to make a copy before you can edit. Make sure you make only private copies.)

1. [PCA] (10 pts) *Note: You are expected to work out the entire question by hand, and not use any libraries/packages. Plots can be drawn manually/through drawing tools, but the labels and lines in the plots have to figured out manually.*

   Bob wants to transmit the following set of four two-dimensional coordinates to his friend.

   $$X = \begin{bmatrix} 4 & 1 \\ 2 & 3 \\ 5 & 4 \\ 1 & 0 \end{bmatrix}$$

   However, due to transmission bandwidth, he is contrained to send only four one-dimensional coordinates. He learns about the PCA algorithm and wants to apply it.

   (a) [6 pts] Find the unit-vector principal components of X. Given that Bob is constrained to send four 1-D coordinates, which principal component would you suggest him to pick, and why? Show your work.

   (b) [4 pts] Bob plots the coordinates of X as in Figure 1. To obtain the 1-D transformation of his 2-D coordinates, he sketches the direction of the principal component

and projects the four 2-D coordinates on this principle component. Show how this plot would look like. Label each of the projected points along with the value of the principal coordinate (note that these labels should be 1-D points).

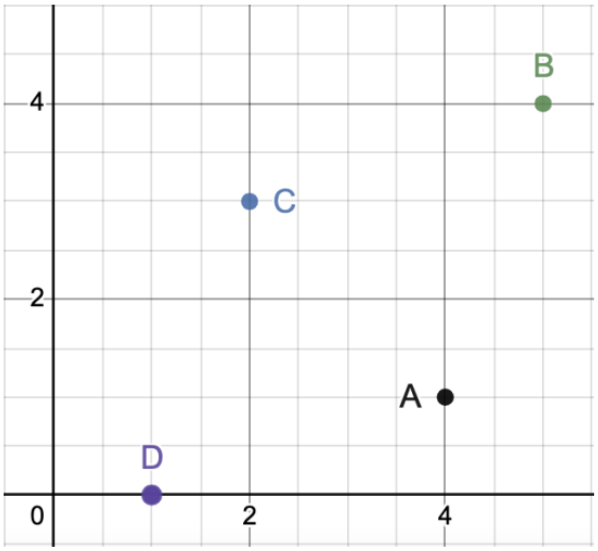

Figure 1: Bob's plot

2. (8 pts) [k-Means] Work through the K-Means clustering algorithm for a dataset with 4 samples, with $K = 2$, and using the L_2 distance. The samples in the dataset are: $A = (2, 3), B = (4, 6), C = (5, 1)$, and $D = (10, 12)$. The initial centroids are chosen as: $(6, 9)$ for cluster 1 and $(8, 4)$ for cluster 2. Recall that in each iteration of K-Means, two things happen: first, cluster assignments are updated, and second, cluster centroids are updated. Work through two such iterations. Report results for each iteration as:

   - A: $d(A, 1)$, $d(A, 2)$
   - B: $d(B, 1)$, $d(B, 2)$
   - C: $d(C, 1)$, $d(C, 2)$
   - D: $d(D, 1)$, $d(D, 2)$
   - cluster 1 members: A, B, etc.
   - cluster 1 updated centroid: (x, y)
   - cluster 2 members: A, B, etc.
   - cluster 2 updated centroid: (x, y)

   where $d(S, c)$ is the L_2 distance from sample S to the cluster c centroid.

3. [Image Filtering/Convolution] (6 pts) We discussed the use of convolution filters for images briefly in class, mostly in the context of CNN. However, before CNNs became popular, convolution filters had already been an essential part of signal processing and

computational photography. You will probably be surprised by how many features in Photoshop or Lightroom can be easily implemented with the correct choice of convolution filter(s). In this question, we will take a look at a few common types of convolution filters for images, and visualize how they would transform the original image. Let's take the following image for example. This is a gray-scale image, where each pixel can be represented by a value between $[0, 1]$, where 0 is black and 1 is white. The gray-scale image itself can be represented by a matrix of shape $width * height$, and we are going to apply 3*3 convolution filters to the matrix. Assume the bias parameter is set to 0 for all these convolution filters.

For the following sub-questions, you will be given a convolution filter, and a few transformed image. Your task is to pick the one that corresponds to the given filter. Your TAs have created a skeleton colab notebook, where you can implement + test out these filters with different images. https://colab.research.google.com/drive/1_mNak1RCxehWBn8Le3LGGHCfbwA2xUTg?usp=sharing.

(a) [2 pts] Consider the following filter:

$$X = \begin{bmatrix} 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \\ 1/9 & 1/9 & 1/9 \end{bmatrix}$$

If we apply X as a convolution filter to the original image, which of the following transformed image will we see? **Briefly justify your choice using one or two sentences.**

- A. 
- C. 

- B. 
- D. 

(b) [2 pts] Now consider this following filter:

$$X = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

If we apply X as a convolution filter to the original image, which of the following transformed image will we see? **Briefly justify your choice using one or two sentences.**

- A. 
- B. 
- C. 
- D. 

(c) [2 pts] Now let's look at a more challenging example:

$$X = \begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

If we apply X as a convolution filter to the original image, which of the following transformed image will we see? **Briefly justify your choice using one or two sentences.**

- A. 
- B. 
- C. 
- D. 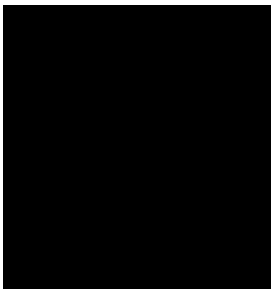

4. [CNNs] (8 pts) Consider the following 3-layer neural network, with layers enumerated starting from the input.

   (a) **Conv2d**: Input channels = 3, Output channels = 5, Kernel size = 5×5, Stride = 1, Padding = 0

   (b) **ReLU**

   (c) **MaxPool2d**: Kernel size = 2×2, Stride = 2, Padding = 0

   (d) **Conv2d**: Input channels = 5, Output channels = 10, Kernel size = 3×3, Stride = 1, Padding = 0

   (e) **ReLU**

   (f) **MaxPool2d**: Kernel size = 2×2, Stride = 2, Padding = 0

   (g) **Conv2d**: Input channels = 10, Output channels = 20, Kernel size = 3×3, Stride = 1, Padding = 0

   (h) **ReLU**

   (i) **MaxPool2d**: Kernel size = 2×2, Stride = 2, Padding = 0

   If the input image has dimensions (Height (dimension 1): 232, Width (dimension 2): 232, Channels (dimension 3): 3), compute the following:

(a) Output height (dimension 1) after all layers

(b) Output width (dimension 2) after all layers

(c) Output number of channels (dimension 3) after all layers

(d) Total number of learnable parameters in the network

2 Python Programming Questions

A IPython notebook is linked on the class website. It will tell you everything you need to do, and provide starter code. Remember to include the plots and answer the questions in your written homework submission!