

This homework is due on Sunday, April 17, 2022 at 11:59PM. Self-grades and HW Resubmissions are due the following Sunday, April 24, 2022 at 11:59PM.

NOTE: The last two problems on this homework discuss concepts that will be covered in the lecture on Tuesday, April 12. As such, we are extending the deadline for this homework to give you extra time to work on those problems. All future homeworks will run on a "Sunday due date, following Sunday resubmission" schedule as well.

1. Reading Lecture Notes

Staying up to date with lectures is an important part of the learning process in this course. Here are links to the notes that you need to read for this week: [Note 15](#) [Note 16](#) [Note 17](#)

- (a) If the outer product form of the SVD of A is $\sum_{i=1}^r \sigma_i \vec{u}_i \vec{v}_i^T$, **what is the best rank k approximation for A , which we'll call X , that minimizes the Frobenius norm $\|A - X\|_F$?** Assume that $k \leq r$.
- (b) We have a data matrix X and its SVD $U\Sigma V^T$. When performing PCA, **what are the principal components if the columns of X are our data points? What are the principal components if the rows of X are our data points?**

2. Rank 1 Decomposition and Image Compression With SVD

In this problem, we will introduce an important application of the singular value decomposition (SVD) called low-rank approximation in the context of image compression. Given a matrix $A \in \mathbb{R}^{m \times n}$ of rank $r \leq \min\{m, n\}$, we saw in [Note 16](#) that we can write A using the outer-product form of the SVD:

$$A = \sum_{i=1}^r \sigma_i \vec{u}_i \vec{v}_i^\top. \quad (1)$$

If our matrix is high-rank, i.e., $r \approx \min\{m, n\}$, then almost all the σ_i will be nonzero and non-negligible. However, if the data has some linear, low-rank essential structure, as is usually the case with real data such as images, most of our singular values will be very small (but usually nonzero due to noise or disturbances). If, say, the data has intrinsic linear rank ℓ , then the first ℓ singular values are large, and the remaining $r - \ell$ are small:

$$A = \sum_{i=1}^r \sigma_i \vec{u}_i \vec{v}_i^\top \quad (2)$$

$$= \sum_{i=1}^{\ell} \sigma_i \vec{u}_i \vec{v}_i^\top + \sum_{i=\ell+1}^r \underbrace{\sigma_i}_{\approx 0} \vec{u}_i \vec{v}_i^\top \quad (3)$$

$$\approx \sum_{i=1}^{\ell} \sigma_i \vec{u}_i \vec{v}_i^\top. \quad (4)$$

This motivates approximating the data as


$$A_\ell := \sum_{i=1}^{\ell} \sigma_i \vec{u}_i \vec{v}_i^\top \quad (5)$$

and using this compressed data for further analysis. As images are readily represented as matrices, in this problem, we will see how to make use of this low rank approximation technique to compress images without severely deteriorating the image quality.

Throughout this problem we will work in the `image_compression.ipynb` Jupyter Notebook. Part 1 of the notebook relates to finding rank 1 decompositions of images, while Parts 2 and 3 focus on image compression. Make sure to read through all the sections of the notebook to understand how we will use what we know about the SVD in order to compress images! You should write all your answers in the written submission and you don't need to submit the notebook.

We will start by decomposing a few images into linear combinations of rank 1 matrices. Remember that outer product of two vectors $\vec{s} \vec{g}^\top$ gives a rank 1 matrix. It has rank 1 because the column span is one-dimensional — multiples of \vec{s} only — and the row span is also one dimensional — multiples of \vec{g}^\top only. More generally, this exercise should provide intuition into how the outer product form of the SVD decomposes the underlying matrix into a sum of constituent rank 1 matrices with different weightings given by the singular values.

- (a) Read and run through the cells in Part 1. Here we first consider a standard 4×4 checkerboard shown in [Figure 1](#). Assume that black colors represent -1 and that white colors represent 1 .

Figure 1: 4×4 checkerboard.

In particular, the checkerboard is given by the following 4×4 matrix L :

$$L = \begin{bmatrix} 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \\ 1 & -1 & 1 & -1 \\ -1 & 1 & -1 & 1 \end{bmatrix} \quad (6)$$

Express L as a linear combination of outer products.

(*HINT: In order to determine how many rank 1 matrices you need to combine to represent the matrix, find the rank of the matrix you are trying to represent.*)

- (b) We next consider the same checkerboard given by the following 4×4 matrix C where black colors now represent 0 and that white colors represent 1.

$$C = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix} \quad (7)$$

Express C as a linear combination of outer products. (*HINT: Note that C is a rank 2 matrix and therefore you need 2 outer products to represent it. To find these individual outer products consider the columns of C .*)

- (c) Now consider the Swiss flag shown in Figure 2. Assume that red colors represent 0 and that white colors represent 1.




Figure 2: Swiss flag.

Assume that the Swiss flag is given by the following 5×5 matrix S :

$$S = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (8)$$

Express S as a linear combination of outer products.

- (d) Run all the cells in Part 2. Thoroughly read through the `rank_k_approx` function. **Since we are now only using the top k singular values, in terms of k , m , and n , how many real numbers do we require to describe the rank k approximation of A ?**
- (e) Run the cells in Part 3 of the notebook. Notice there are two images: One compressed, one uncompressed. **Can you easily visually detect any difference between the two images even though we are using only about 40% of the data?**
- (f) Play around with the slider for the kitten picture. **What is the lowest acceptable value of k that you can go without losing too much image quality? What were the memory savings?**
NOTE: Use your best judgment since this depends on eyesight, screen resolution, etc.
- (g) Plot the singular values for the image of the US flag. **What do you notice about the singular values here in comparison to the kitten's singular values? What does this mean for our low rank compression?**
- (h) Play around with the interactive slider for the US flag. **What is the lowest acceptable value of k here? What is the memory saving?**

3. Frobenius Norm

In this problem we will investigate the basic properties of the Frobenius norm.

Similar to how the norm of vector $\vec{x} \in \mathbb{R}^n$ is defined as $\|\vec{x}\| = \sqrt{\sum_{i=1}^n x_i^2}$, the Frobenius norm of a matrix $A \in \mathbb{R}^{m \times n}$ is defined as

$$\|A\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |A_{ij}|^2}. \quad (9)$$

A_{ij} is the entry in the i^{th} row and the j^{th} column. This is basically the norm that comes from treating a matrix like a big vector filled with numbers.

(a) With the above definitions, **show that for a 2×2 matrix A :**

$$\|A\|_F = \sqrt{\text{tr}(A^T A)}. \quad (10)$$

Note: The trace of a matrix is the sum of its diagonal entries. For example, let $A \in \mathbb{R}^{m \times n}$, then,

$$\text{tr}(A) = \sum_{i=1}^{\min(n,m)} A_{ii} \quad (11)$$

Think about how/whether this expression eq. (10) generalizes to general $m \times n$ matrices.

(b) **Show that if U and V are square orthonormal matrices, then**

$$\|UA\|_F = \|AV\|_F = \|A\|_F. \quad (12)$$

(HINT: Use the trace interpretation from part (a).)

(c) **Use the SVD decomposition to show that $\|A\|_F = \sqrt{\sum_{i=1}^n \sigma_i^2}$, where $\sigma_1, \dots, \sigma_n$ are the singular values of A .**

(HINT: The previous part might be quite useful.)

4. Movie Ratings and PCA

Recall from the lecture on PCA that we can think of movie ratings as a structured set of data. For every person i and movie j , we have that person's rating $R_{i,j}$ (thought of as a real number).

Suppose that there are m movies and n people. Let's think about arranging this data into a big $n \times m$ matrix R with people corresponding to rows and movies corresponding to columns. So the entry in the i -th row and j -th column should be $R_{i,j}$ above. Note that this is organized differently from how it was in lecture. Each row corresponds to a unique person and each column to a unique movie.

$$R = \begin{bmatrix} R_{11} & R_{12} & \cdots & R_{1m} \\ R_{21} & R_{22} & \cdots & R_{2m} \\ \vdots & \vdots & \cdots & \vdots \\ R_{n1} & R_{n2} & \cdots & R_{nm} \end{bmatrix} \quad (13)$$

- (a) Suppose we believe that there is actually an underlying pattern to this rating data and that a separate study has revealed that every movie is characterized by a set of characteristics: say action and comedy. This means that every movie j has a pair of numbers a_j (for action) and c_j (for comedy) that define it. At the same time, every person i has a pair of sensitivities f_i and g_i that defines that person's preferences for action vs. comedy movies respectively. A person i will rate the movie j as $R_{i,j} = f_i a_j + g_i c_j$.

If we arrange the sensitivities into a pair of n -dimensional vectors \vec{f}, \vec{g} for our group of n people, and the movie characteristics into a pair of m -dimensional vectors \vec{a}, \vec{c} for our group of m movies, **use outer products to express the rating matrix R in terms of these vectors $\vec{f}, \vec{g}, \vec{a}, \vec{c}$.**

- (b) Now suppose that we want to discover the underlying nature of movies from the data R itself. Suppose for this part, that you have four observed rating data vectors (corresponding to four different movies being rated by six individuals).

All of the movie data vectors just happened to be multiples of the following 6-dimensional vector

$$\vec{w} = \begin{bmatrix} 2 \\ -2 \\ 3 \\ -4 \\ 4 \\ 0 \end{bmatrix}. \quad (\text{For your convenience, note that } \|\vec{w}\| = 7.)$$

You arrange the movie data vectors as the columns of a matrix R given by:

$$R = \begin{bmatrix} | & | & | & | \\ -\vec{w} & -2\vec{w} & 2\vec{w} & 4\vec{w} \\ | & | & | & | \end{bmatrix} \quad (14)$$

You want to perform PCA (for movies) using the SVD of the matrix R to better understand the pattern in your data.

The first "principal component vector" is a unit vector that tells which direction we would want to project the columns of R onto to get the best rank-1 approximation for R .

Find this first principal component vector of the columns of R to explain the nature of your movie data points.

(HINT: You don't need to actually compute any SVDs in this simple case. Also, be sure to think about what size vector you want as the answer. Don't forget that you want a unit vector!)

- (c) Suppose that now, we have two more data points (corresponding to two more movies being rated by the same set of six people, i.e. we added two columns to our matrix) that are multiples of a different vector \vec{p} where:

$$\vec{p} = \begin{bmatrix} 6 \\ 3 \\ -2 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \text{ (For your convenience, note that } \|\vec{p}\| = 7 \text{ and that } \vec{p}^\top \vec{w} = 0.)$$

We augment our ratings data matrix with these two new data points to get:

$$R = \begin{bmatrix} | & | & | & | & | & | \\ -\vec{w} & -2\vec{w} & 2\vec{w} & 4\vec{w} & -3\vec{p} & 3\vec{p} \\ | & | & | & | & | & | \end{bmatrix} \quad (15)$$

Find the first two principal components corresponding to the nonzero singular values of R . This is what we would use to best project the movie data points onto a two-dimensional subspace.


What is the first principal component vector? What is the second principal component vector? Justify your answer. (HINT: Think about the inner product of \vec{w} and \vec{p} and what that implies for being able to appropriately decompose R . Again, very little computation is required here.)

(d) In the previous part, you had

$$R = \begin{bmatrix} | & | & | & | & | & | \\ -\vec{w} & -2\vec{w} & 2\vec{w} & 4\vec{w} & -3\vec{p} & 3\vec{p} \\ | & | & | & | & | & | \end{bmatrix}$$

with $\|\vec{w}\| = 7$ and $\|\vec{p}\| = 7$, satisfying $\vec{p}^\top \vec{w} = 0$.

If we use \vec{r}_i to denote the i -th column of R , **plot the movie data points \vec{r}_i (for all i) projected onto the first and second principal component vectors along the columns of R .** The coordinate along the first principal component should be represented by horizontal axis and the coordinate along the second principal component should be the vertical axis. **Label each point, and the axes. Remember that principal component vectors are normalized.**



5. Brain-Machine Interface and Neuron Sorting

The iPython notebook `pca_BMI.ipynb` will guide you through the process of analyzing brain machine interface data leveraging the SVD. This will help you to prepare for the project.

For SIXT33N (your robot car), you will need to use the SVD and PCA to classify your sound inputs so that your car does what you tell it to do.

Brain-Machine interfaces (BMIs) are a way for a brain to directly communicate to an external device. They're often used for research, mapping, assisting or repairing human cognitive or sensory-motor function.

Data was collected that shows waveform traces of (simulated) brain waves of a subject whose arm is pointing in certain directions over time. These waveform traces are gathered from electrodes inserted into the brain, but the electrodes are generally recording more than 1 neuron at the same time — the brain is crowded after all. To make predictions based on neuron firing rates, we need to be able to distinguish waveforms that come from different neurons.

Each neuron has its own waveform “signature” shape unique to that neuron. This is due to the physical characteristics of neurons, such as their physical shape and structure. It is impossible to know beforehand how many neurons an electrode measures or what each neuron’s waveform looks like, so we must first separate the waveforms from the different neurons near the electrode.

The goal of this problem is to see how we can use the SVD (in its PCA manifestation) and clustering to decide which neuron fired. We will first look at a case where there are only two neurons, then a case where there are three neurons. The neurons have also been presorted using a professional software package, so we can check our model against presorted data.

Please complete the notebook by following the instructions given.

Task 1: Two Neuron Waveform Sorting

- The data you have are traces of length 32 timesteps. The data is arranged along the rows of the matrix, each row is one trace. Import the data sets and see the average waveform for each presorted neuron by running the corresponding cells in the `.ipynb`. **What is the shape of the training data matrix `three_neurons_training`?**
- What do the columns and rows of the training data matrix: `three_neurons_training` represent?**

We can approximate each waveform in a lower dimensional space using PCA. In your implementation of PCA you will decide how to choose these components.

- You will be using the SVD in your PCA function. **What matrices does the SVD return? What are the dimensions of these matrices for the SVD of `three_neurons_training`?**
- To represent each waveform in a lower dimensional space we want a basis for the time signals of the waveforms of the neurons. To construct this basis we start by taking the SVD (of `three_neurons_training`). **Which of the U and V matrices can we use to construct our desired basis?**
- Read through `PCA_train` and implement `PCA_project`.** We will use these functions throughout the rest of the notebook, so make sure you understand them.
- Call `PCA_train` on `two_neurons_training` and plot the 2 principal components.** Note that since the dataset is randomized, you might get different plots every time you run the second code cell of this notebook (the `_make_training_set` function). Note that these components have no real “physical” relationship to the actual shape of the neuron plots. They are a part of a basis, not the representation of exemplar traces.

- (g) Before we project onto our principal components, let us project our data onto 2 random 32 length vectors. **Run the corresponding part in the .ipynb file and comment on the separation of the 2 neuron's distinctive trace shapes in this projected basis.** Do you think that it would be easy to tell them apart just by looking at their projection into these two random dimensions?
- (h) **Project `two_neurons_test` onto the two principal components found using the SVD and produce a 2D scatter plot in the new basis.** We will also try projecting the presorted data containing 2 neurons so we can see how the model behaves on the 2 neurons. **Comment on the difference between the projections here and part (g) where you projected onto random directions.**
- (i) Now we will repeat this process for three principal components. **Do you see a significant improvement with 3 principal components?**

Task 2: Three Neuron Sorting

- (j) In the previous part we sorted 2 neurons using two or three principal components. In this part we are now dealing with 3 neurons. **Replicate what we did in the previous parts by finding and projecting our data on the first two principal components of this three neuron dataset in the .ipynb.**
- (k) Now call `PCA_train` on `three_neurons_training` and **plot the 3 principal components. Do the same classification process as the 2 neuron data, but now with the 3 neuron data. Compare your model's behavior with that of the presorted data.** The `plot_3D` function will be useful to view the results.
- (l) **How many principal components do you actually need to cluster the 3 neurons?** (*HINT: Think about whether there was an increase in separability between the last two parts.*)
- (m) Now that we know how to project our data onto a basis where it is clearly separable we can classify our points and determine which neuron fired. (This is what we need to know if we want to use the firings of different neurons to build a BMI.)
To classify our points we use the following algorithm:
 - i. Find centroids: points that represent the average waveform of a particular neuron firing, in the basis of principal components.
 - ii. Classify each incoming point by assigning it the value of the centroid closest to it (i.e., declare that the neuron waveform that a point represents is the same as the neuron waveform of the centroid the point is closest to).

Since we already have some presorted data, we can use this information to aid in classification. We can calculate our centroids by taking the empirical mean of the presorted data corresponding to a particular neuron firing. **Use this method to find centroids in the corresponding section of the .ipynb file. Then use `which_neuron` on the `two_classified` data set and count the number of times each neuron fired.**

6. (OPTIONAL) Make Your Own Problem.

Write your own problem about content covered in the course thus far, and provide a thorough solution to it.

NOTE: This can be a totally new problem, a modification on an existing problem, or a Jupyter part for a problem that previously didn't have one. Please cite all sources for anything (including course material) that you used as inspiration.

NOTE: High-quality problems may be used as inspiration for the problems we choose to put on future homeworks or exams.

7. Homework Process and Study Group

Citing sources and collaborators are an important part of life, including being a student!

We also want to understand what resources you find helpful and how much time homework is taking, so we can change things in the future if possible.

- (a) **What sources (if any) did you use as you worked through the homework?**
- (b) **If you worked with someone on this homework, who did you work with?**
List names and student ID's. (In case of homework party, you can also just describe the group.)
- (c) **Roughly how many total hours did you work on this homework? Write it down here where you'll need to remember it for the self-grade form.**

Contributors:

- Daniel Abraham.
- Ashwin Vangipuram.
- Tanmay Gautam.
- Aditya Arun.
- Kouros Hakhmaneshi.
- Antroy Roy Chowdhury.
- Nikhil Shinde.
- Nathan Lambert.
- Anant Sahai.
- Wahid Rahman.