

Assignment #4

Due: 12:00 noon EDT, Oct 5, 2020

Upload at: <https://www.gradescope.com/courses/135869/assignments/687666>

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 to, say, 95% or so to get everything on the page.

Problem 1 (13pts)

For the following problems, you may use any combination of hand-calculation and Python code (as long as you include the code you used).

(A) Compute an orthonormal basis of the *kernel* of

$$A = \begin{bmatrix} 1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & 1 & 1 \end{bmatrix}$$

(B) Write down an orthonormal basis for the *image* of A .

Problem 2 (30pts)

You've encountered power series (e.g., Taylor series) before in other classes, but one thing you may not have realized is that you can construct *matrix functions* from *matrix power series*. That is, if you have a function $f(\cdot)$ that has a convergent power series representation:

$$f(x) = \sum_{i=0}^{\infty} a_i x^i$$

then you can generally write a similar matrix version for square symmetric matrices \mathbf{X} using the same a_i :

$$F(\mathbf{X}) = \sum_{i=0}^{\infty} a_i \mathbf{X}^i$$

- (A) The matrix version $F(\mathbf{X})$ turns out to just apply the scalar $f(x)$ to each eigenvalue independently. Explain why. (Hint: How would a diagonalized version of \mathbf{X} interact with the power series?)
- (B) In power series there is a notion of **radius of convergence**. How would you expect this concept to generalize to square symmetric matrices?
- (C) One important example is where the function $f(x)$ is the exponential function. I can take any square symmetric matrix, and if I compute its matrix exponential, I get a positive definite matrix. Explain why.
- (D) These kinds of matrix functions lead to some interesting computational tricks. For example: if I have a positive definite matrix \mathbf{A} and I take the *exponential* of the *trace* of the *matrix logarithm* (assuming it exists), what quantity have I computed?

Problem 3 (25pts)

In general, computing the determinant of an $n \times n$ matrix scales as n^3 in computational cost. When the matrix is highly structured, however, it can sometimes be possible to take advantage of that structure to save computation for quantities such as the determinant. One example of such structure is in *tridiagonal* matrices, which look like this:

$$T = \begin{bmatrix} a_1 & b_1 & 0 & 0 & \cdots & 0 \\ c_1 & a_2 & b_2 & 0 & & 0 \\ 0 & c_2 & a_3 & \ddots & & \vdots \\ 0 & 0 & \ddots & \ddots & b_{n-2} & 0 \\ \vdots & & & c_{n-2} & a_{n-1} & b_{n-1} \\ 0 & 0 & \cdots & 0 & c_{n-1} & a_n \end{bmatrix}$$

Such matrices can come up when simulating a physical system with local structure, e.g., a spring-mass system.

- (A) We would like to compute the determinant of the matrix T above, which we are assuming is invertible. Let d_m denote the determinant of the upper left $m \times m$ submatrix. So, $d_1 = a_1$ and $d_n = \det(T)$ (i.e., computing the whole determinant). Use **expansion by minors** on the $m + 1$ st row to compute d_{m+1} in terms of d_m , d_{m-1} , and any of the a_i , b_i , or c_i . To avoid special cases, take $d_0 = 1$ and $d_i = 0$ for $i < 0$. Carefully show how you arrived at your answer.
- (B) This recurrence relation motivates a recursive algorithm for computing $\det(T)$ by computing d_1, d_2, \dots, d_n in order. How would you expect the computational cost to scale for computing the determinant of a tridiagonal matrix?

Problem 4 (30pts)

This problem relates to content that will be presented the week of Sep 28. If you wish to get started before then, look at sections 4.5-4.6 and 10.1-10.6 in your book (though here we use eigenvectors and not SVD).

One of the most important techniques in data analysis is **principal component analysis** or PCA. PCA tries to find a way to represent high-dimensional data in a low-dimensional way so that human brains can reason about it. It tries to identify the “important” directions in a data set and represent the data just in that basis. PCA does this by computing the empirical covariance matrix of the data (we’ll learn more about that in a couple of weeks), and then looking at the eigenvectors of it that correspond to the largest eigenvalues.

- (A) Load `mnist2000.pkl` into a Colab notebook. Take the $2000 \times 28 \times 28$ tensor of training data and reshape it so that it is a 2000×784 matrix, where the rows are “unrolled” image vectors. Typically in PCA, one first centers the data. Center the data by subtracting off the mean image; you did a very similar procedure in HW2.
- (B) Now compute the “scatter matrix” which is the 784×784 matrix you get from multiplying data matrix by its transpose, making sure that you get it so the data dimension is the one being summed over.
- (C) This scatter matrix is square and symmetric, so use the `eigh` function in the `numpy.linalg` package to compute the eigenvalues and eigenvectors. Plot the eigenvalues in decreasing order.
- (D) Read the documentation for `eigh` and figure out how to get the “big” eigenvectors. For each of the top five eigenvectors, reshape them into 28×28 images and use `imshow` to render them.
- (E) Now, create a low-dimensional representation of the data. Take the 2000×784 matrix and multiply it by each of the top two eigenvectors. This takes all 2000 data, each of which are 784-dimensional, and gives them two-dimensional coordinates. Make a scatter plot of these two-dimensional coordinates.
- (F) That scatter plot doesn’t really give you much of a visualization. Here’s some starter code to build a more interesting figure. It takes the two-dimensional projection and builds a “scatter plot” where the images themselves are rendered instead of dots. Here I have the projections in a 2000×2 matrix called `proj`, which I modify so that all the values are in `[0, 1]`.

```
# Make the projections into [0,1]
proj = proj - np.min(proj, axis=0)
proj = proj / np.max(proj, axis=0)

# Create a 12" x 12" figure.
viz_fig = plt.figure(figsize=(12.,12.))

# Get the figure width and height in pixels.
width, height = viz_fig.get_size_inches()*viz_fig.dpi

plt.plot() # Colab seems to require this to render.

# Loop over images. Could do all 2000 but it's crowded.
for ii in range(400):
    # Render each image in a location on the figure.
    plt.figure(figsize=(12.,12.))
    plt.imshow(train_images[ii,:,:],
               xo=proj[ii,1]*width,
               yo=(proj[ii,0]*height-150), # hack to make visible
               origin='upper')
```

Modify this code to work with your projections and make a visualization of the MNIST digits. Do you see any interesting structure?

Problem 5 (2pts)

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

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

Changelog

- 20 September 2020 – updated for fall 2020.
- 24 February 2020 – Clarified problem two.
- 23 February 2020 – Initial version.