

Instructor: A.A. Ahmadi

TAs: Beneventano, Chaudhry, Hua, Li, Lok

Due on Thursday, September 28, 2023, at 1:30pm EST, on Gradescope

For all problems that involve coding, please include your code. Our instructions are given with MATLAB in mind (recommended language), but you are free to use a different language.


Problem 1: Image compression and remembering John Horton Conway¹

Figure 1: The image to be compressed.

If we were to rank the topics in computational linear algebra in terms of how frequently they arise in applications, *singular value decomposition* (SVD) would probably take one of the very top spots. In this problem, we introduce you to this concept and present one of its applications in image processing. The SVD also has applications in many other domains, notably in statistics.

Let A be a real $m \times n$ matrix of rank r . (Recall that the rank of A is the number of linearly independent columns of A .) The singular value decomposition of A is a decomposition of the form

$$A = U\Sigma V^T,$$

¹John Conway (1937–2020) was an English mathematician and a Professor of Mathematics at Princeton University. He is known for his fundamental contributions to numerous areas of mathematics. You may enjoy reading about Conway [here](#), or playing his [Game of Life here](#).

where U, Σ , and V are respectively $m \times m, m \times n$, and $n \times n$; U and V are orthogonal matrices (i.e., satisfy $U^T U = I$ and $V^T V = I$), and Σ is a matrix with r positive scalars $\sigma_1, \dots, \sigma_r$ on the diagonal of its upper left $r \times r$ block and zeros everywhere else. The scalars $\sigma_1, \dots, \sigma_r$ are called the *singular values* of A . They are given as

$$\sigma_i = \sqrt{i\text{-th eigenvalue of } A^T A},$$

and by convention they appear in descending order:

$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_r.$$

The columns of U and V are respectively called the left and right *singular vectors* of A and can be obtained by taking an orthonormal set of eigenvectors² for the matrices AA^T and $A^T A$. In Matlab, the command `svd` handles these eigenvector computations for you and outputs the three matrices U, Σ, V .

First, you need to answer some very basic questions about the SVD.

1. Show that eigenvalues of $A^T A$ are always nonnegative. (Hence singular values are well-defined as real, nonnegative scalars.)
2. Show that if u_i and u_j are eigenvectors of $A^T A$ associated to distinct eigenvalues λ_i, λ_j , then u_i and u_j are orthogonal.

What does this have to do with image compression? Let A be a real $m \times n$ matrix with an SVD given by $A = U\Sigma V^T$ defined as above. For a positive integer $k \leq \min\{m, n\}$, we let $A_{(k)}$ denote an $m \times n$ matrix which is an ‘‘approximation’’ of the matrix A obtained from its top k singular values and singular vectors. Formally, we have

$$A_{(k)} := U_{(k)} \Sigma_{(k)} V_{(k)}^T,$$

where $U_{(k)}$ has the first k columns of U , $V_{(k)}$ has the first k columns of V , and $\Sigma_{(k)}$ is the upper left $k \times k$ block of Σ . A well-known result in optimization states that $A_{(k)}$ is a minimizer of the optimization problem

$$\min_{B \in \mathbb{R}^{m \times n}, \text{rank}(B) \leq k} \|A - B\|_F.$$

Here, $\|\cdot\|_F$ denotes the so-called Frobenius norm of a matrix defined as $\|C\|_F = \sqrt{\sum_{i,j} C_{i,j}^2}$.

²By an orthonormal set of vectors we mean a collection of vectors that are pairwise orthogonal and each have 2-norm equal to one.

In words, the result states that among all $m \times n$ matrices of rank at most k , the matrix $A_{(k)}$ obtained from truncating the SVD is the one that best approximates A . The benefit in approximating a matrix with low-rank matrices is that low-rank matrices admit a much more succinct representation. Let's see how this applies to our image compression problem.

Download the file `conway.jpg` into your Matlab path. You can read this in by typing:

```
1 A=imread('conway.jpg');
2 A=im2double(A);
3 A=rgb2gray(A);
```

If you are using Python, you may wish to use the package `pillow` and `numpy`:

```
1 from PIL import Image
2 import numpy as np
3 image = Image.open('conway.jpg')
4 A = np.array(image, dtype=float) / 255.0
5 A = np.dot(A, [0.2989, 0.5870, 0.1140])
```

The result is a 1120×960 matrix A , with each entry representing a single pixel in the picture with a number between 0 and 1. To upload this picture on Instagram, you would need to upload $1120 \times 960 = 1075200$ numbers (pixels).

3. For $k = 40, 80, 120, 160$, use Matlab to compute $A_{(k)}$ as defined above. Report the value of $\|A - A_{(k)}\|_F$ in each case. (Include your code for this part and the next.)
4. Use the commands `subplot` and `imshow` to produce on the same figure the original image, as well as your compressed images $A_{(k)}$ for $k = 40, 80, 120, 160$. Label your subplots. In addition, produce two separate plots demonstrating (i) $\|A - A_{(k)}\|_F$ versus k , and (ii) "total savings" versus k . Total savings is to be interpreted as the answer to the question: How many fewer numbers do you need in order to store $A_{(k)}$ than you did to store A . Explain why this number is equal to $mn - (n + m + 1)k$. How much are you saving for $k = 160$?
5. Use the Matlab function `imwrite` and upload the two images `imshow(A)` and `imshow(A_{(160)})` side by side on Instagram with the caption `#PrincetonORF363`. Ask your friends if they can tell which image is the original and which is the compressed version. (We are joking. There's no grade assigned to this part of the problem!)

Problem 2: Find all the local minimizers, local maximizers, global minimizers, and global maximizers of the following function over \mathbb{R}^2 (or argue if some do not exist):

$$f(x_1, x_2) = \frac{1}{2}x_1^2 + 4x_1x_2 + \frac{1}{2}x_2^2 - x_2^3.$$

Plot this function using Matlab and the command `ezsurf` and rotate the surface around with your mouse. Does the number of local minima and maxima that you found agree with what you see on the plot?

Problem 3: Properties of positive semidefinite matrices

Prove or disprove the following statements. All matrices are $n \times n$, symmetric, and real.

- (a) Suppose $A \succeq 0$. Then the largest entry in absolute value of A must be on the diagonal.³
- (b) If $A \succeq 0$ and $\text{trace}(A) = 0$, then $A = 0$.
- (c) If $A \succeq 0, B \succeq 0$, and $A + B = 0$, then $A = B = 0$.
- (d) If $A \succeq 0, B \succeq 0$, and $AB = 0$, then $A = 0$ or $B = 0$.

Problem 4: Copositive matrices

An $n \times n$ real symmetric matrix Q is said to be *copositive* if $x^T Q x \geq 0$ for all $x \in \mathbb{R}^n$ such that $x \geq 0$. (The inequality on x is elementwise.)

- (a) Prove that the set of $n \times n$ copositive matrices is convex. Show that the set of $n \times n$ noncopositive matrices is nonconvex unless $n = 1$.
- (b) Give an example of a matrix that is copositive but neither positive semidefinite nor elementwise nonnegative. (You have to prove all claims about the example that you produce.)

³In other words, the largest entry of $|A|$ must be on the diagonal, where $|A|$ is the matrix whose entries are the absolute values of those of A .