Exercises from Strang P.358 4,9,10,15,16,24,32,36; P.370 1,3,6,7; P. 379 2,4,6,8,9,15,18,23,26; P.398 1,2,4,15,17

Extra Question: For the two flags in Question 3 from Page 370, use the below to compute their SVDs.

Recall that if $\vec{u}_1, \ldots, \vec{u}_r$ are the left-singular vectors of A, $\vec{v}_1, \ldots, \vec{v}_r$ are the right-singular vectors of A, and $\sigma_1, \ldots, \sigma_r$ are the singular values of A, then the SVD gives A as a sum of outer products:

$$A = \sigma_1 \vec{u}_1 \vec{v}_1^T + \ldots + \sigma_r \vec{u}_r \vec{v}_r^T.$$

If the latter σ's are very small, their corresponding outer products contribute very little to A. So perhaps we can drop them without losing information. To do so, we just drop the last columns of U and V, and make Σ a smaller square.

Reconstruct each of the flags using just the first singular vectors and values (i.e. $k = 1$ below). By considering the relative sizes of the singular values for each, explain your results as they compare to the original matrices. What does $k = 2$ give you? Explain.

SVD in Sage

```
A=matrix(RDF,[[1,2,3],[4,5,6]])
U,S,Vt=A.SVD()
```

(Note that SAGE outputs U, Σ, and V^T. That is, the V is already transposed.)

Test:

```
print(A)
print(U*S*Vt)
```

RDF stands for 'Real Double Field', one of the ways Sage represents real numbers as floating point (specifically, double precision). Sage cannot do exact SVD (yet?).

Low Dimension Reconstruction To reconstruct a matrix from just the first k singular vectors and values:

```
k=2
print(U[:,:k]*S[:k,:k]*V[:,:k].T)
```