12 Matlab interlude: Orthogonal Matrices, and the SVD

In this class we stepped back from our study of finite difference methods to revisit topics on orthogonal matrices, and low-rank approximation using the SVD, with applications to image processing.

The notes are available at
www.maths.nuigalway.ie/~niall/MA519/lab1/lab1.pdf

When I get the chance, I'll flesh out this page with some details about low-rank approximation.

In the mean time, here is the black and white version of the official NUI Galway logo. It is represented as a 154×500 pixel image in 616000 bytes (8 bytes per pixel).

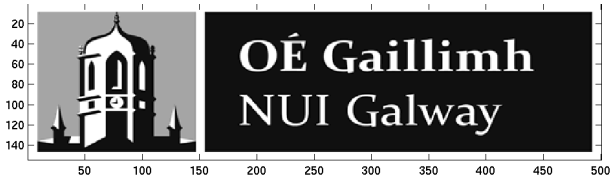

Figure 12.1: Original image of the NUI Galway logo. It has rank 125

.

Now here in Figure 12.2 is the distribution of the singular values. We can see from this that the matrix does not have full rank. In fact, the rank is 125.

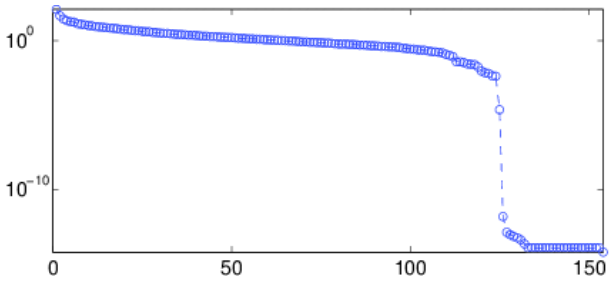

Figure 12.2: The distribution of the singular values of the image shown in Figure 12.1

Here in Figure 12.3 is the rank 2 approximation. It required just 1310 bytes, but is not have a good likeness to the original. Can you tell from looking at it that it is a rank 2 image?

In Figure 12.4 we show the rank 10 approximation. It required 6550 bytes, but the image is quite clear.

Finally, in Figure 12.5 I give that rank 40 approximation (26200 bytes). It is difficult to discerne any difference between this a the one in Figure 12.1.

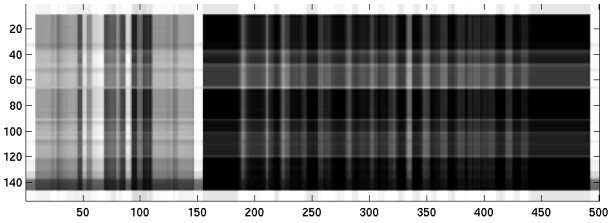

Figure 12.3: Rank 2

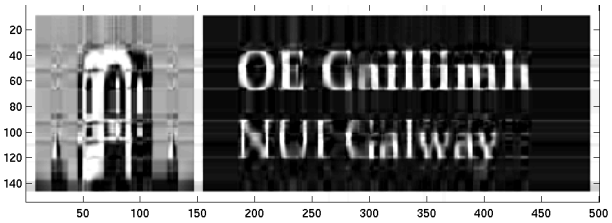

Figure 12.4: Rank 10 approximation of the image in Figure 12.1

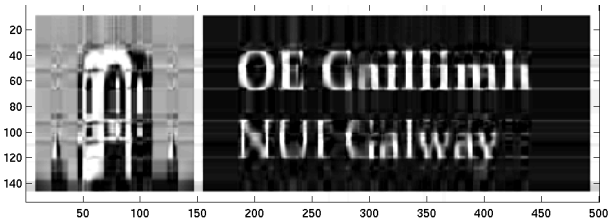

Figure 12.5: Rank 40 approximation of the image in Figure 12.1