The UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL

Comp 665 Imaging, Graphics, and Vision
Fall 2008

Problem Set #4
Issued Tuesday, 11/18/08; Due Tuesday, 12/2/08

Homework Information: Some of the problems are probably too long to be done the night before the due date, so plan accordingly. Late homework will penalized as explained on the course Web site. Feel free to get help from others, but the work you hand in should be your own.

Problem 1. "Tensor Decomposition"

Using the 144×144×3 Mandrill image construct a 48×48×27 image stack by de-interleaving every third pixel on every third scan line. This is accomplished with the following few lines of Python:

```
data = asarray(Image.open("MandrillTiny.png"))
H, W, D = data.shape
tensor = empty((H/3,W/3,D*9))
for i in xrange(9):
    tensor[:,:,3*i:3*(i+1)] = data[(i/3)%3:H:3,(i%3):W:3,:]
```

Perform a tensor decomposition of the rows, columns, and tubes of this stack.

A) How many row vectors are necessary to represent 95% of the variance under the assumptions of this model? Likewise, how many column and tube vectors are necessary to represent 95% of the variance?
B) Repeat part A for 98% of the total variance.
C) Solve for the core tensor using the row, column, and tube dimensions from Part A, and print its entries (Hint: It has less than 30 entries).
D) What is the Mean Absolute Error achieved using the sum of multilinear components implied by the core tensor in part C?
E) Compare the number of PCA coefficients needed to achieve the same 95% of the total variance as the tensor decomposition in part A, when each of the 48 × 48 subimages are treated as a vector. How do the total sizes of the required row, column, and tube Eigenvectors compare to the sizes of the PCA factors?

Problem 2. "Which Wavelet?"

In the following questions the term *detail coefficients* refers to all wavelet-coefficient matrix entries other than the mean value, $c_{0,0}$.

A) Perform a *standard* 2D wavelet transform of the 256×256 gray-scale image, MandrillGray.pgn, by applying a Haar wavelet transform to each row, and then to each column. Print an image of the wavelet-coefficient matrix (Make sure that you consider that many coefficients are negative). What are the magnitudes of the largest 10 detail coefficients?
B) Apply the following mapping to all detail coefficients of the LH, HL, and HH quadrants ($c_{i,j}$ where $i \geq 128$ or $j \geq 128$).

$$c'_{i,j} = \begin{cases} 10 & \textit{if } c_{i,j} > 8 \\ -10 & \textit{if } c_{i,j} < -8 \\ 0 & \textit{otherwise} \end{cases}$$

Reconstruct the modified image, compute its MAE, and print it.

C) Next, we consider a *non-standard* 2D Haar wavelet transform of the same image. Perform one-level of averaging and differencing to the pixel values of each row, then apply one-level of averaging and differencing to each column. Complete the transformation by repeating this operation recursively to the quadrant containing both averages, until the average is computed for the entire image. Print an image of the resulting wavelet-coefficient matrix (It should be noticeably different than the one from part A). Once again compute the magnitudes of the largest 10 detail coefficients?

D) Apply the mapping function described in part B to the non-standard transform's detail coefficients of quadrants HL, LH, and HH. Reconstruct the modified image, compute its MAE, and print it.

E) Compare and contrast the two 2D wavelet-transform approaches. Consider any anomalies that you may have noticed from the given exercises. Discuss the computational complexity of the two approaches.