Lab Three: Image Compression

Linear Algebra

College of the Atlantic

Due Friday, May 17, 2024

Instructions: Complete the following exercises in groups of two students; submit a single report for your group. If you want, you can write directly on the handout. Be sure to include complete explanations of the work you have done and justification for your conclusions. Scan your work and submit on google classroom. Please make sure that both lab partners' names are on the assignment.

Use https://cocalc.com for sage as needed.

These problems are based on exercises written by Matt Boelkins, which are based on problems from Understanding Linear Algebra, by David Austin.

In the JPEG compression algorithm, we are interested in representing a digital image using the smallest amount of data possible. By converting from the RGB color model to the YC_bC_r color model, we are concentrating the most visually important data into a single quantity. This is helpful because we can safely ignore some of the data in the chrominance values since that data is not as visually important. In many cases, we can effectively store an image and re-render it accurately using less than 10% of the storage space the original image would require.

The JPEG compression algorithm begins by breaking an image (which might be 1440×1468 pixels or larger) into 8×8 blocks of pixels. Each pixel is represented as a color using three coordinates in the luminance-chrominance color model (YC_bC_r). For instance, here are the luminance values in one particular 8×8 block from a certain image.

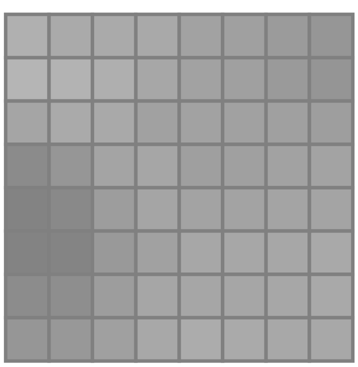

176	170	170	169	162	160	155	150
181	179	175	167	162	160	154	149
165	170	169	161	162	161	160	158
139	150	164	166	159	160	162	163
131	137	157	165	163	163	164	164
131	132	153	161	167	167	167	169
140	142	157	166	166	166	167	169
150	152	160	168	172	170	168	168

There are 64 values here. The JPEG compression algorithm works by finding a way of representing this information using only a small fraction of the 64 values. Our strategy is to perform a change of basis to take advantage of the fact that the luminance values do not change significantly over the block. Rather than recording the luminance of each of the pixels individually, this change of basis will allow us to record the average luminance along with some information about how the individual values vary from the average.

In class we looked at the first column of this matrix. For this exercise, we'll look at the **right-most column**:

$$\vec{x} = \begin{bmatrix} 150 \\ 149 \\ 158 \\ 163 \\ 164 \\ 169 \\ 169 \\ 168 \end{bmatrix} \tag{1}$$

As we did in class, we'll represent the vector $\vec{x}$ in a new basis $\mathcal{F}$, given by:

$$\vec{v}_0 = \begin{bmatrix} \cos\left(\frac{(2\cdot0+1)0\pi}{16}\right) \\ \cos\left(\frac{(2\cdot1+1)0\pi}{16}\right) \\ \cos\left(\frac{(2\cdot2+1)0\pi}{16}\right) \\ \vdots \\ \cos\left(\frac{(2\cdot7+1)0\pi}{16}\right) \end{bmatrix}, \vec{v}_1 = \begin{bmatrix} \cos\left(\frac{(2\cdot0+1)1\pi}{16}\right) \\ \cos\left(\frac{(2\cdot1+1)1\pi}{16}\right) \\ \cos\left(\frac{(2\cdot2+1)1\pi}{16}\right) \\ \vdots \\ \cos\left(\frac{(2\cdot7+1)1\pi}{16}\right) \end{bmatrix}, \ldots, \vec{v}_7 = \begin{bmatrix} \cos\left(\frac{(2\cdot0+1)7\pi}{16}\right) \\ \cos\left(\frac{(2\cdot1+1)7\pi}{16}\right) \\ \cos\left(\frac{(2\cdot2+1)7\pi}{16}\right) \\ \vdots \\ \cos\left(\frac{(2\cdot7+1)7\pi}{16}\right) \end{bmatrix} \tag{2}$$

The process of expressing $\vec{x}$ in this new basis is known as the *discrete Fourier transform.*

Because $\mathcal{F}$ is a basis for $\mathbb{R}^8$, we can represent any vector $\vec{x}$ in standard coordinates instead in its $C-$coordinates. When we do so, we will denote the coordinates of $\vec{x}$ relative to C with the notation

$$\{\vec{x}\}_{\mathcal{F}} = \vec{f} = \begin{bmatrix} F_0 \\ F_1 \\ F_2 \\ F_3 \\ F_4 \\ F_5 \\ F_6 \\ F_7 \end{bmatrix},$$

where components of $\vec{f}$ are called the *Fourier coefficients* of the vector $\vec{x}$. Said differently, we can use the Fourier coefficients to write

$$\vec{x} = F_0\vec{v}_0 + F_1\vec{v}_1 + \cdots + F_7\vec{v}_7.$$

1. The matrix P transforms a vector from the Fourier basis to the standard basis. That is,

$$P\vec{f} = \vec{x}\,. \tag{3}$$

   Explain how you would form the matrix P.

2. If you were given a vector $\vec{x}$, explain how would you convert it into the Fourier basis

3. Find the Fourier representation of the vector $\vec{x}$ of Eq. (1).

4. Explain why F_0, the first component of $\vec{f}$, is the average of the components of $\vec{x}$.

5. Explain why it is not surprising that F_7 is almost zero.

6. Next, form an approximation $\vec{f}_{\text{approx}}$ to $\vec{f}$ by rounding the Fourier coefficients to the nearest integer.

7. Now find the vector $\vec{x}_{\text{approx}}$ by converting from the $\mathcal{F}$ basis back to the standard basis.

8. Compare $\vec{x}_{\text{approx}}$ to the original vector $\vec{x}$ by computing $|\vec{x} - \vec{x}_{\text{approx}}|$. What is the largest difference in the components of $\vec{x}_{\text{approx}}$ and $\vec{x}$?

9. Express this difference as a percent of the original $\vec{x}$ value.

10. What percentage of the components of $\vec{f}$ get turned into zeros when forming $\vec{f}_{\text{approx}}$?

11. Explain briefly how the process of forming $\vec{f}_{\text{approx}}$ is a form of image compression.