

ECE438 - Laboratory 10: Image Processing (Week 2)

October 18, 2008

1 Introduction

This is the second part of a two week experiment in image processing. In the [first week](#), we covered the fundamentals of digital monochrome images, intensity histograms, pointwise transformations, gamma correction, and image enhancement based on filtering.

During this week, we will cover some fundamental concepts of *color* images. This will include a brief description on how humans perceive color, followed by descriptions of two standard *color spaces*. We will also discuss an application known as *halftoning*, which is the process of converting a gray scale image into a binary image.

2 Color Images

2.1 Background on Color

Color is a perceptual phenomenon related to the human response to different wavelengths of light, mainly in the region of 400 to 700 nanometers (nm). The perception of color arises from the sensitivities of three types of neurochemical sensors in the retina, known as the *long* (L), *medium* (M), and *short* (S) *cones*. The response of these sensors to photons is shown in Figure 1. Note that each sensor responds to a range of wavelengths.

Due to this property of the human visual system, all colors can be modeled as combinations of the three *primary color* components: red (R), green (G), and blue (B). For the purpose of standardization, the CIE (Commission International de l'Eclairage — the International Commission on Illumination) designated the following wavelength values for the three primary colors: blue = $435.8nm$, green = $546.1nm$, and red = $700nm$.

The relative amounts of the three primary colors of light required to produce a color of a given wavelength are called *tristimulus values*. Figure 2 shows the plot of tristimulus values using the CIE primary colors. Notice that some of the tristimulus values are *negative*, which indicates that colors at those wavelengths cannot be reproduced by the CIE primary colors.

Questions or comments concerning this laboratory should be directed to Prof. Charles A. Bouman, School of Electrical and Computer Engineering, Purdue University, West Lafayette IN 47907; (765) 494-0340; bouman@ecn.purdue.edu




Figure 1: Relative photon sensitivity of long (L), medium (M), and short (S) cones.




Figure 2: Plot of tristimulus values using CIE primary colors.

2.2 Color Spaces

A *color space* allows us to represent all the colors perceived by human beings. We previously noted that weighted combinations of stimuli at three wavelengths are sufficient to describe all the colors we perceive. These wavelengths form a natural basis, or coordinate system, from which the color measurement process can be described. In this lab, we will examine two common color spaces: RGB and YC_bC_r . For more information, refer to [1].

- RGB space is one of the most popular color spaces, and is based on the tristimulus theory of human vision, as described above. The RGB space is a hardware-oriented model, and is thus primarily used in computer monitors and other raster devices. Based upon this color space, each pixel of a digital color image has three components: red, green, and blue.
- YC_bC_r space is another important color space model. This is a gamma corrected space defined by the CCIR (International Radio Consultative Committee), and is mainly used in the digital video paradigm. This space consists of *luminance* (Y) and *chrominance* (C_bC_r) components. The importance of the YC_bC_r space comes from the fact that the

human visual system perceives a color stimulus in terms of luminance and chrominance attributes, rather than in terms of R , G , and B values. The relation between YC_bC_r space and gamma corrected RGB space is given by the following linear transformation.

$$\begin{aligned} Y &= 0.299R + 0.587G + 0.114B \\ C_b &= 0.564(B - Y) + 128 \\ C_r &= 0.713(R - Y) + 128 \end{aligned} \quad (1)$$

In YC_bC_r , the luminance parameter is related to an overall intensity of the image. The chrominance components are a measure of the relative intensities of the blue and red components. The inverse of the transformation in equation (1) can easily be shown to be the following.

$$\begin{aligned} R &= Y + 1.4025(C_r - 128) \\ G &= Y - 0.3443(C_b - 128) - 0.7144 * (C_r - 128) \\ B &= Y + 1.7730(C_b - 128) \end{aligned} \quad (2)$$

2.3 Color Exercise



You will be displaying both color and monochrome images in the following exercises. Matlab's *image* command can be used for both image types, but care must be taken for the command to work properly. Please see the [help on the image command](#) for details.

Download the RGB color image file [girl.tif](#), and load it into Matlab using the *imread* command. Check the size of the Matlab array for this image by typing *whos*. Notice that this is a three dimensional array of type *uint8*. It contains three gray scale image planes corresponding to the red, green, and blue components for each pixel. Since each color pixel is represented by three bytes, this is commonly known as a 24-bit image. Display the color image using

```
image(A);
axis('image');
```

where A is the 3-D RGB array.

You can extract each of the color components using the following commands.

```
RGB = imread('girl.tif'); % color image is loaded into matrix RGB
```

```

R = RGB(:,:,1);           % extract red component from RGB
G = RGB(:,:,2);           % extract green component from RGB
B = RGB(:,:,3);           % extract blue component from RGB

```

Use the *subplot* and *image* commands to plot the original image, along with each of the three color components. Note that while the original is a color image, each color component separately is a monochrome image. Use the syntax `subplot(2,2,n)`, where $n = 1, 2, 3, 4$, to place the four images in the same figure. Place a title on each of the images, and print the figure (use a color printer).

We will now examine the YC_bC_r color space representation. Download the file [ybcbr.mat](#), and load it into Matlab using `load ybcbr`. This file contains a Matlab array for a color image in YC_bC_r format. The array contains three gray scale image planes that correspond to the *luminance* (Y) and two *chrominance* (C_bC_r) components. Use `subplot(3,1,n)` and *image* to display each of the components in the same figure. Place a title on each of the three monochrome images, and print the figure.

In order to properly display this color image, we need to convert it to RGB format. Write a Matlab function that will perform the transformation of equation (2). It should accept a 3-D YC_bC_r image array as input, and return a 3-D RGB image array.

Now, convert the *ybcbr* array to an RGB representation and display the color image. Remember to convert the result to type *uint8* before using the *image* command.

An interesting property of the human visual system, with respect to the YC_bC_r color space, is that we are much more sensitive to distortion in the luminance component than in the chrominance components. To illustrate this, we will smooth each of these components with a Gaussian filter and view the results.

You may have noticed when you loaded *ybcbr.mat* into Matlab that you also loaded a 5×5 matrix, *h*. This is a 5×5 Gaussian filter with $\sigma^2 = 2.0$. (See the first week of the experiment for more details on this type of filter.) Alter the *ybcbr* array by filtering only the luminance component, `ybcbr(:,:,1)`, using the Gaussian filter (use the *filter2* function). Convert the result to RGB , and display it using *image*. Now alter *ybcbr* by filtering both chrominance components, `ybcbr(:,:,2)` and `ybcbr(:,:,3)`, using the Gaussian filter. Convert this result to RGB , and display it using *image*.

Use `subplot(3,1,n)` to place the original and two filtered versions of the *ybcbr* image in the same figure. Place a title on each of the images, and print the figure (in color). Do you see a significant difference between the filtered versions and the original image? This is the reason that YC_bC_r is often used for digital video. Since we are not very sensitive to corruption of the chrominance components, we can afford to lose some information in the encoding process.

INLAB REPORT:

1. Submit the figure containing the components of *girl.tif*.
2. Submit the figure containing the components of *ycbcr*.
3. Submit your code for the transformation from YC_bC_r to RGB .
4. Submit the figure containing the original and filtered versions of *ycbcr*. Comment on the result of filtering the luminance and chrominance components of this image. Based on this, what conclusion can you draw about the human visual system?

3 Halftoning

In this section, we will cover a useful image processing technique called *halftoning*. The process of halftoning is required in many present day electronic applications such as facsimile (FAX), electronic scanning/copying, laser printing, and low bandwidth remote sensing.

3.1 Binary Images

As was discussed in the first week of this lab, an 8-bit monochrome image allows 256 distinct gray levels. Such images can be displayed on a computer monitor if the hardware supports the required number intensity levels. However, some output devices print or display images with much fewer gray levels. In the extreme case, the gray scale images must be converted to binary images, where pixels can only be black or white.



Figure 3: (a) Original gray scale image. (b) Binary image produced by simple fixed thresholding.

The simplest way of converting to a binary image is based on *thresholding*, i.e. two-level (one-bit) quantization. Let $f(i, j)$ be a gray scale image, and $b(i, j)$ be the corresponding

binary image based on thresholding. For a given threshold T , the binary image is computed as the following:

$$b(i, j) = \begin{cases} 255 & \text{if } f(i, j) > T \\ 0 & \text{else} \end{cases} \quad (3)$$

Figure 3 shows an example of conversion to a binary image via thresholding, using $T = 80$.

It can be seen in Figure 3 that the binary image is not “shaded” properly—an artifact known as *false contouring*. False contouring occurs when quantizing at low bit rates (one bit in this case) because the quantization error is dependent upon the input signal. If one reduces this dependence, the visual quality of the binary image is usually enhanced.

One method of reducing the signal dependence of the quantization error is to add uniformly distributed white noise to the input image prior to quantization. To each pixel of the gray scale image $f(i, j)$, a white random number n in the range $[-A, A]$ is added, and then the resulting image is quantized by a one-bit quantizer, as in equation (3). The result of this method is illustrated in Figure 4, where the additive noise is uniform over $[-40, 40]$. Notice that even though the resulting binary image is somewhat noisy, the false contouring has been significantly reduced.



Figure 4: Random noise binarization.

3.2 Ordered Dithering

Halftone images are binary images that appear to have a gray scale rendition. Although the random thresholding technique described in Sec. 3.1 can be used to produce a halftone image, it is not often used in real applications since it yields very noisy results. In this section, we will describe a better halftoning technique known as *ordered dithering*.

The human visual system tends to average a region around a pixel instead of treating each pixel individually, thus it is possible to create the illusion of many gray levels in a binary image, even though there are actually only two gray levels. With 2×2 binary pixel grids,

we can represent 5 different “effective” intensity levels, as shown in Figure 5. Similarly for 3×3 grids, we can represent 10 distinct gray levels. In dithering, we replace blocks of the original image with these types of binary grid patterns.




Figure 5: Five different patterns of 2×2 binary pixel grids.

Remember from Sec. 3.1 that false contouring artifacts can be reduced if we can reduce the signal dependence of the quantization error. We showed that adding uniform noise to the monochrome image can be used to achieve this decorrelation. An alternative method would be to use a variable threshold value for the quantization process.

Ordered dithering consists of comparing blocks of the original image to a 2-D grid, known as a *dither pattern*. Each element of the block is then quantized using the corresponding value in the dither matrix as a threshold. The values in the dither matrix are fixed, but are typically different from each other. Because the threshold value varies between adjacent pixels, some decorrelation from the quantization error is achieved, which has the effect of reducing false contouring.

The following is an example of a 2×2 dither matrix,

$$T(i, j) = 255 * \begin{bmatrix} 5/8 & 3/8 \\ 1/8 & 7/8 \end{bmatrix} \quad (4)$$

This is a part of a general class of optimum dither patterns known as *Bayer matrices*. The values of the threshold matrix $T(i, j)$ are determined by the order that pixels turn “ON”. The order can be put in the form of an *index matrix*. For a Bayer matrix of size 2, the index matrix $I(i, j)$ is given by

$$I(i, j) = \begin{bmatrix} 3 & 2 \\ 1 & 4 \end{bmatrix} \quad (5)$$

and the relation between $T(i, j)$ and $I(i, j)$ is given by

$$T(i, j) = 255(I(i, j) - 0.5)/n^2 \quad (6)$$

where n^2 is the total number of elements in the matrix.

Figure 6 shows the halftone image produced by Bayer dithering of size 4. It is clear from the figure that the halftone image provides good detail rendition. However the inherent square grid patterns are visible in the halftone image.




Figure 6: The halftone image produced by Bayer dithering of size 4.

3.3 Error Diffusion

Another method for halftoning is random dithering by *error diffusion*. In this case, the pixels are quantized in a specific order (raster ordering¹ is commonly used), and the residual quantization error for the current pixel is propagated (diffused) forward to unquantized pixels. This keeps the overall intensity of the output binary image closer to the input gray scale intensity.




Figure 7: Block diagram of the error diffusion method.

Figure 7 is a block diagram that illustrates the method of error diffusion. The current input pixel $f(i, j)$ is modified by means of past quantization errors to give a modified input $\tilde{f}(i, j)$. This pixel is then quantized to a binary value by Q , using some threshold T . The error $e(i, j)$ is defined as

$$e(i, j) = \tilde{f}(i, j) - b(i, j) \quad (7)$$

where $b(i, j)$ is the quantized binary image.

¹Raster ordering of an image orients the pixels from left to right, and then top to bottom. This is similar to the order that a CRT scans the electron beam across the screen.

The error $e(i, j)$ of quantizing the current pixel is diffused to “future” pixels by means of a two-dimensional weighting filter $h(i, j)$, known as the *diffusion filter*. The process of modifying an input pixel by past errors can be represented by the following recursive relationship.

$$\tilde{f}(i, j) = f(i, j) + \sum_{k,l \in S} h(k, l)e(i - k, j - l) \quad (8)$$

The most popular error diffusion method, proposed by Floyd and Steinberg, uses the diffusion filter shown in Figure 8. Since the filter coefficients sum to one, the local average value of the quantized image is equal to the local average gray scale value. Figure 9 shows the halftone image produced by Floyd and Steinberg error diffusion. Compared to the ordered dither halftoning, the error diffusion method can be seen to have better contrast performance. However, it can be seen in Figure 9 that error diffusion tends to create “streaking” artifacts, known as *worm* patterns.

		•	7/16
3/16	5/16	1/16	

Figure 8: The error diffusion filter proposed by Floyd and Steinberg.




Figure 9: A halftone image produced by the Floyd and Steinberg error diffusion method.

3.4 Halftoning Exercise

Download [house.tif](#)

An important preliminary note: If your display software (e.g Matlab) resizes your image before rendering, a halftoned image will probably not be rendered properly. For example, a subsampling filter will result in gray pixels in the displayed image! To prevent this in Matlab, use the *truesize* command just after the *image* command. This will assign one monitor pixel for each image pixel.

We will now implement the halftoning techniques described above. **Save** the result of each method in MAT files so that you may later analyze and compare their performance. Download the image file [house.tif](#) and read it into Matlab. Print out a copy of this image.

First try the simple thresholding technique based on equation (3), using $T = 108$, and display the result. In Matlab, an easy way to threshold an image X is to use the command $Y = 255*(X>T);$. Label the quantized image, and print it out.

Now create an “absolute error” image by subtracting the binary from the original image, and then taking the absolute value. The degree to which the original image is present in the error image is a measure of signal dependence of the quantization error. Label and print out the error image.

Compute the mean square error (MSE), which is defined by

$$MSE = \frac{1}{NM} \sum_{i,j} \{f(i,j) - b(i,j)\}^2 \quad (9)$$

where NM is the total number of pixels in each image. Note the MSE on the printout of the quantized image.

Now try implementing Bayer dithering of size 4. You will first have to compute the dither pattern. The index matrix for a dither pattern of size 4 is given by

$$I(i,j) = \begin{bmatrix} 12 & 8 & 10 & 6 \\ 4 & 16 & 2 & 14 \\ 9 & 5 & 11 & 7 \\ 1 & 13 & 3 & 15 \end{bmatrix} \quad (10)$$

Based on this index matrix and equation (6), create the corresponding threshold matrix.

For ordered dithering, it is easiest to perform the thresholding of the image all at once. This can be done by creating a large threshold matrix by repeating the 4×4 dither pattern. For example, the command $T = [T \ T; \ T \ T];$ will increase the dimensions of T by 2. If this is repeated until T is at least as large as the original image, T can then be trimmed so that it is the same size as the image. The thresholding can then be performed using the command $Y = 255*(X>T);$.

As above, compute an error image and calculate the MSE. Print out the quantized image, the error image, and note the MSE.

Now try halftoning via the error diffusion technique, using a threshold $T = 108$ and the diffusion filter in Figure 8. It is most straightforward to implement this by performing the following steps on each pixel in raster order:

1. Initialize an output image matrix with zeros.
2. Quantize the current pixel using using the threshold T , and place the result in the output matrix.
3. Compute the quantization error by subtracting the binary pixel from the gray scale pixel.
4. Add scaled versions of this error to “future” pixels of the original image, as depicted by the diffusion filter of Figure 8.
5. Move on to the next pixel.

You do not have to quantize the outer border of the image.

As above, compute an error image and calculate the MSE. Print out the quantized image, the error image, and note the MSE.

The human visual system naturally lowpass filters halftone images. To analyze this phenomenon, filter each of the halftone images with the Gaussian lowpass filter h that you loaded in the previous section (from *ycbcr.mat*), and measure the MSE of the filtered versions. Make a table that contains the MSE’s for both filtered and nonfiltered halftone images for each of the three methods. Does lowpass filtering reduce the MSE for each method?

INLAB REPORT:

1. Hand in the original image and the three binary images. Make sure that they are all labeled, and that the mean square errors are noted on the binary images.
2. Compare the performance of the three methods based on the visual quality of the halftoned images. Also compare the resultant MSE’s. Is the MSE consistent with the visual quality?
3. Submit the three error images. Which method appears to be the least signal dependent? Does the signal dependence seem to be correlated with the visual quality?
4. Compare the MSE’s of the filtered versions with the nonfiltered versions for each method. What is the implication of these observations with respect to how we perceive halftone images.

References

- [1] J.M. Kasson and W. Plouffe, “An analysis of selected computer interchange color spaces,” *ACM Trans. Graphics*, vol. 11, no. 4, pp. 373–405, 1992.