INTRODUCTION TO IMAGE PROCESSING

We are surrounded with digital images – on computer screens and HDTVs and in digital photography and video. Unlike traditional film cameras, when we take a picture with a digital camera we are able to easily manipulate the image by altering its appearance electronically if we so desire. Think about some digital images you have worked with in the past. What are you able to do to your picture? Can you think of a few other examples of where image processing is used?

If we treat a digital image as an "input signal" and the resulting image as the output then we can use image processing to describe the image alteration process. Figure 1 below shows a simple block diagram of image processing. How would you describe the result of the "image processing block" based on the input and output images below?

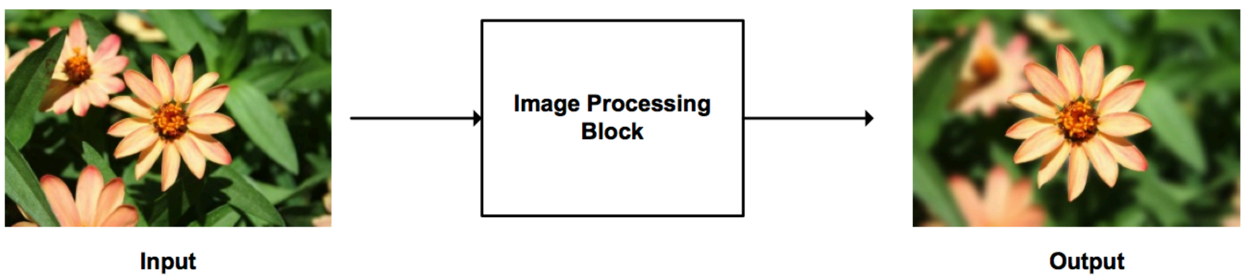

Figure 1: Image processing on an image captured with a digital camera

In this lab we will learn about some of the theory and techniques that go on inside of the image processing block. To do this we will have to draw on our understanding of both physics and mathematics.

What is a digital image?

A digital image is defined as a two-dimensional representation (ie. a picture) of some subject, usually a physical object or a person. The image is composed of many small dots known as picture elements or pixels. Because the dots are so small and close together, our brain is able to assemble to dots into a meaningful image. However, if we were to enlarge the picture we would eventually see each pixel and would lose focus of the overall image. Oftentimes, the original object is difficult to identify in these "pixelated" versions.

The number of pixels being used to represent an image determines the amount of detail in the image. This is known as resolution. For example, older "standard definition" TVs have a resolution of 525 pixels (height) by 700 pixels (width), while today's 1080p HDTVs have a resolution of 1080 pixels by 1920 pixels. That means that there are 367,500 pixels for standard TVs and 993,600 pixels for 1080p HDTVs – there are more than 600,000 more pixels crammed into the HDTVs. Think about why the "packedness" of the pixels makes for a better image…

How are pixels used to represent the image?

Digital images come in various formats, such as black & white, grayscale or color. In each of these cases, in order to represent the image, each pixel of an image is assigned a number that represents the "intensity of color."

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

Black and White Images

In a black & white image a pixel that is black would be assigned a zero (full intensity of black) and a pixel that is white would be assigned a one (no black). The entire image would therefore be represented by a grid (array) of zeros and ones, where the intensity of each pixel is represented by the appropriate number in the array.

> ***Technically...*** Black and white images are called a one bit image -- Each pixel is described by one binary digit (0 or 1). Figure 2 below is an example of a binary image.

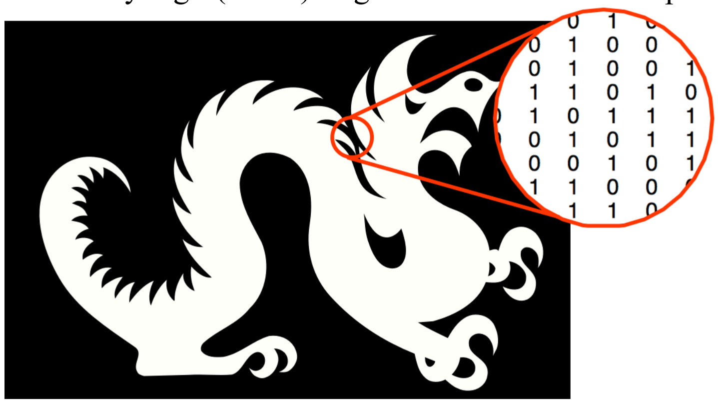

Figure 2: A Black and White (Binary) Image

Grayscale Images

What happens if instead of simply representing each “pixel” of an image with black or white, we allow different “intensities of black” (shades of gray)? How would this image differ from a black and white image? To describe grayscale digital images, we permit light intensities to fall between the range of black (zero) and white (one), we can represent each pixel by a number between zero and one that is proportional to the “intensity” of black. For example, an intensity of 0.2 would be “darker” than an intensity of 0.8. Typically, each pixel is assigned a number between 0 and 255, where 0 represents the complete absence of light (black) and 255 represents complete saturation of light (white). In this case, there are 255 levels of increasing intensity of light.

> ***Technically...*** We need to increase the number of bits that are used to represent the intensity of a pixel. Instead of using only one bit we may now use 8 bits to define the intensity of a pixel. An 8-bit sequence ('11001001') can take on 256 different values. Therefore each pixel can take on 256 intensities if represented by an 8-bit sequence. What color would 0 refer to? What about 255? What about 100?
>
> If we increase the depth of the image from 8-bits to 16-bits we now can represent each pixel by 65,536 different numbers. What can you conclude about the bit depth of an image? What would happen to the file size of the image by increasing the image depth? Figure 3 below shows an image allowing for 256 different intensities for each pixel. How might this process produce numbers that are similar (or equivalent) to the percentages discussed above?

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

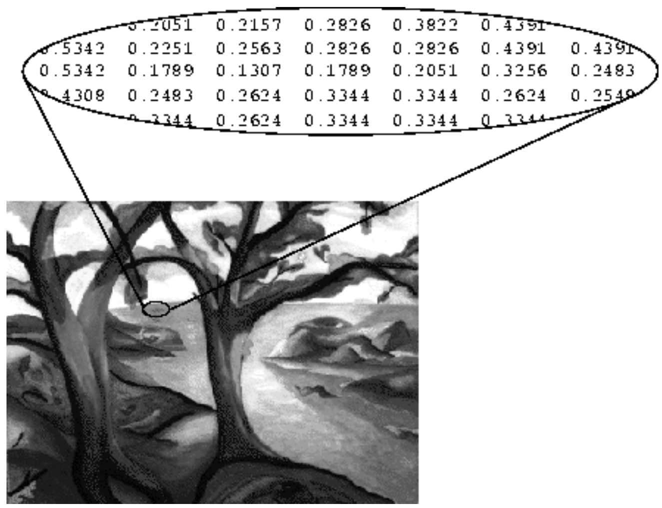

Figure 3: An 8-bit (grayscale) image

Color and Light

Because light also makes color (and color images) possible, a brief overview of light is needed in order to understand and represent color images. The electromagnetic energy spectrum is comprised of energy waves also known as light. Our eyes are capable of detecting only a small portion of light that we call the visible spectrum. This spectrum is bounded by the colors we see as violet and red. At the higher end of the visible spectrum beyond violet is light of shorter wavelengths known as ultraviolet and x-rays. At the lower end of this spectrum beyond red is light of longer wavelengths known as infrared and radio.

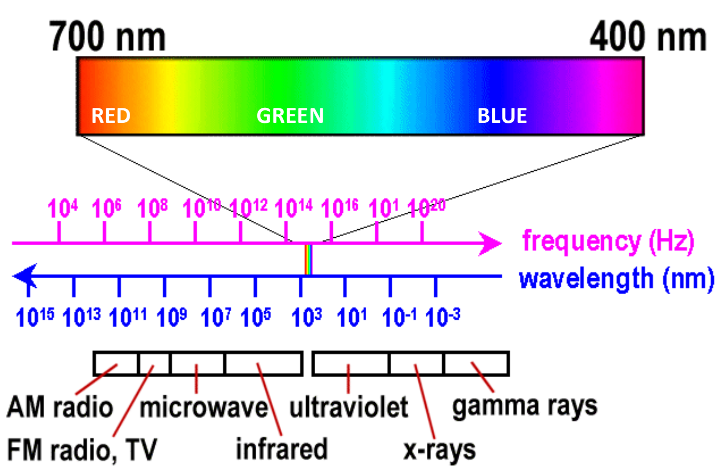

Figure 4: The visible spectrum

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

As shown in figure 4 all colors visible to us exist in between red and violet. If the visible portion of the light spectrum were divided into thirds, the predominant colors would be red, blue and green (RGB). These colors are considered to be the modern primary colors.

The remaining colors can be created by combining red, blue, and green light. For example, when all RGB are all in the same intensity, we see white. When there is the absence of all RGB, we see black. When you combine G and R in equal intensities, we see yellow. Other combinations (and the spectrum of color) is shown below in the color wheel shown in Figure 6.

> ***Technically...*** The primary colors can be thought of as forming a large triangle. The colors in between the primary colors namely, cyan, magenta and yellow (CMY) would form another triangle. Combining two of the additive primary colors produces the secondary colors cyan, magenta, and yellow and a combination of all three produces white.

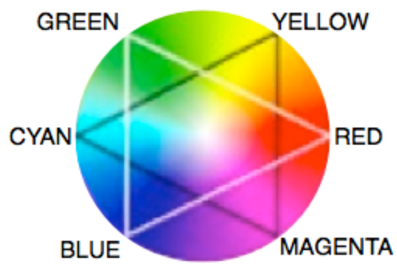

Figure 5: The Color Wheel

Color Digital Images
The LCD monitor on your computer uses three tiny light sources for each pixel. Each light source is either red, blue or green and each can vary the intensity of light it produces. Therefore each pixel can represent any color from black to white by varying the intensity of these three light sources. How would you represent black? White? Cyan?

Previously, for black and white and grayscale images, we only required one array to represent the entire digital image because we were only working with the intensity of black. For color we need to vary the intensity of three colors for each pixel. Therefore we require three different arrays for each image.

An image that is m pixels by n pixels are stored as an m-by-n-by-3 array, where m is the number of rows and n is the number of columns. Each of the three m-by-n "layers" represents the intensity of one of the primary colors for each individual pixel. See Figure 6 for an example of representing a color digital image.

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

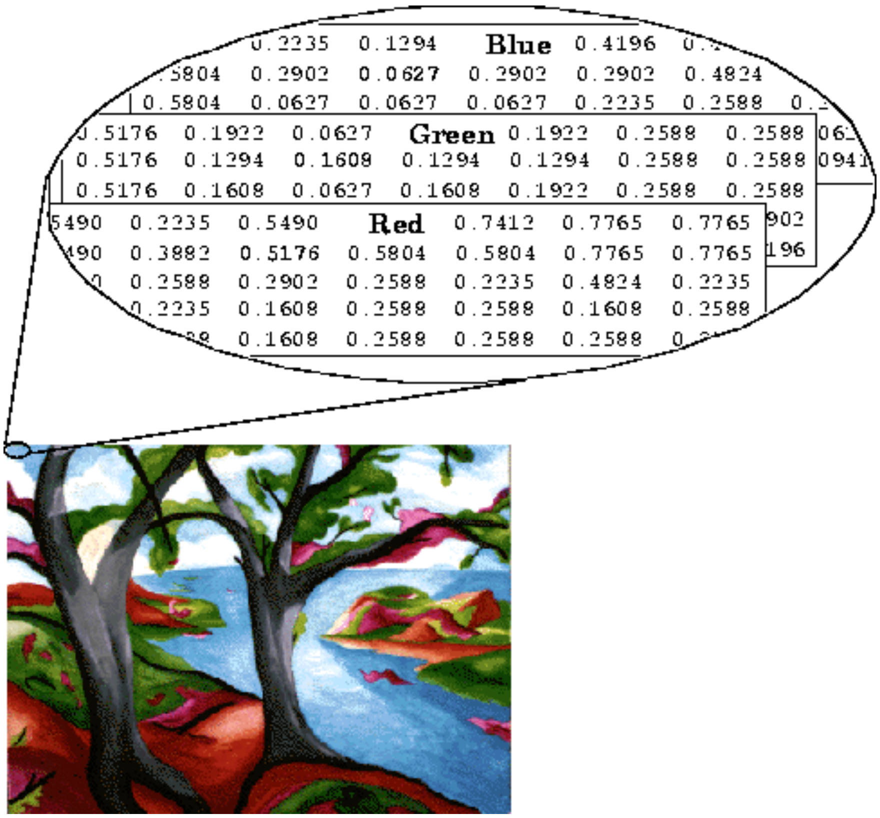

Figure 6: An RGB image

Now that we understand and have built a foundation for the representation of digital images we are ready to begin learning the theory behind a few image processing techniques. Since we have digitized the image by representing it by a sequence of numbers were are able to manipulate it using mathematics on the sequence. Once we have completed our math to alter the 3-dimensional array we can send it to our LCD screen to verify that our image processing performed as expected.

You are now going to work in small groups on a series of activities to familiarize yourselves with digital with the theory behind a few popular image processing techniques.

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

Activity 1: Exploring Pixels on your Computer Screen

In these activities, you will explore the RGB values of the individual pixels that comprise an image on your computer screen. To begin, you will need a digital color meter application. In Apple OSX, open Digital Color Meter (Applications>Utilities>Digital Color Meter); on a PC, you can download "Pixie" at http://www.nattyware.com/pixie.php. With the application open, move the mouse to a part of the screen that is white. Notice the Red, Green and Blue values. Try black next. What are these values? Why? Explore different parts of your screen and the different RGB values for the pixels until you are comfortable with the relationship between the RGB values and the pixel color.

Test yourself: predict the RGB values for various images and icons. Were you close? different colors.

Editing Digital Images: Tinting and Thresholding

Two methods for editing digital images include tinting and the application of a thresholding operator. With tinting, the overall R, G, and B values in an image are adjusted, resulting in an edited image that has an overall different color or "feel." With thresholding, a threshold is chosen and, for each pixel in the image, if a specific value is larger than the threshold, that color is replaced with a zero (i.e. the color is turned off); if the value is smaller than the threshold, it is left unchanged.

Activity 2: Thresholding an Image by Hand

Consider the following grayscale image, which is represented by a 9x9 matrix. Each cell includes the intensity of black, on a proportional scale from black (0 = no light or black) to white (1 = full saturation of light or white).

0.532	0.312	0.364	0.458	0.479	0.371	0.412	0.562	0.411
0.453	0.794	0.651	0.401	0.446	0.497	0.786	0.636	0.443
0.403	0.774	0.675	0.415	0.312	0.462	0.789	0.629	0.479
0.488	0.534	0.512	0.534	0.493	0.570	0.409	0.534	0.493
0.452	0.470	0.531	0.821	0.763	0.662	0.578	0.554	0.583
0.603	0.297	0.361	0.258	0.762	0.417	0.425	0.515	0.707
0.572	0.654	0.584	0.450	0.456	0.443	0.511	0.654	0.427
0.301	0.484	0.899	0.632	0.783	0.611	0.791	0.507	0.413
0.488	0.534	0.512	0.534	0.493	0.570	0.409	0.534	0.493

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).

Thresholding involves selecting a "threshold" value and for every pixel (cell) whose value is greater than the threshold value, turning that color "off" (replacing it with a zero). In the above graph, shade take a threshold of 0.6 and shade all cells that have an intensity greater than 0.6 black. Do you see a hidden image?

Use the blank matrix below to create a hidden image. Share it with your friends… can they figure out your threshold?

Activity 3: Tinting and Thresholding an Image

1. Download the following NASA image of Mars: http://dk12.ece.drexel.edu/MT/mars.jpg and save it to your desktop.
2. Open the "Tinting" applet at http://dk12.ece.drexel.edu/image_guis/Tinting.html or following the link on the main Image Processing Labs page.
3. Click SELECT FILE and then direct the applet to the saved image.
4. Use the sliders to explore what happens if you alter the Red, Green or Blue channels.
    a. Slide R to zero. What do you notice? Can you explain why that happens?
    b. Slide G to zero. What do you notice? Can you explain why that happens?
    c. Slide B to zero. What do you notice? Can you explain why that happens?
    d. Fiddle with the channels to make Mars appear more red. Explain why the slider locations
5. Think about how this activity relates to the previous activity with thresholding. Discuss this with your colleagues.

This material is based upon work supported by the National Science Foundation under Grant No. 0733284. Any opinions, findings and conclusions or recommendations expressed in this websites do not necessarily reflect the views of the National Science Foundation (NSF).