Land Remote Sensing
Pre-lab 2: Basics of Image Processing

Goals for this lab:
→ Introduce ENVI software
→ Learn what a digital image is and how it is displayed
→ Learn the basics of image manipulation
→ Introduce the image histogram and basic contrast stretching

1 What is a Digital Image?

The basic dataset for any and all projects in this class is the digital image. Fortunately, most of us have seen plenty of digital images, either from web pages, email attachments, or any number of other sources. Many of us have even "touched up" digital images with programs like MS Paint or Adobe Photoshop. What makes this possible? We may have told computers to manipulate an image for us, but how does the computer do it? The answer lies with understanding what a digital image is.

A digital image is a grid of picture elements, called *pixels*. Pixels are familiar to most, as when we zoom in on an image on a computer screen enough times the image becomes very blocky, or pixelated. Each pixel represents a certain color, but it is important to realize that pixels are not simply colors—**they are numbers called brightness values or, sometimes, data numbers (DN)**. In other words, a digital image is really a grid or array of numbers stored in a computer file. Whatever program used to display the image turns a gridded number into a color and displays that color on the screen. Pixel values are called brightness values because the higher the number, the brighter the color assigned to them. Because digital images are grids of numbers, we can manipulate them mathematically. This is the basis of image processing.

It's also important to remember what each of the numbers in an image represents. Consider a digital camera in which a grid of detectors is used in place of film. If you were to take a picture with a digital camera, the numbers for each pixel in the image would be proportional to the number of photons that each of the detector cells detected. In a typical remote sensing image, the pixel values represent the number of photons **of a certain wavelength** (or range of wavelengths) that are recorded by the detector. Actually, it is more true to say that most remote sensing images merely start off this way. In most cases, calibrations and corrections are used to convert the photon count images to more useful quantities, e.g. reflectance.

Of course, there are many image files that are made of more than one grid of numbers. Color images, for example, are made of three stacked grids, and it is perfectly appropriate to think of these kinds of image files as being made of three separate images. We call each of the separate images a *band* or a *channel*, which represent the imaged data at a certain wavelength. For a given application in remote sensing, we can be using single-band images (like many radar images), images with several bands (Landsat-7 images have seven bands), or even images with hundreds of bands (AVIRIS images, which have 224 bands). A complete image file, with all bands considered together, is often called an image cube, and example of which is given in Figure 1. In this lab, we will look at some basic processing techniques, starting with single-band images and moving to more complicated multiband images.

2 Single-Band Images

Single-band images are often called *grayscale images* because they are usually not displayed in color (although this is possible). Instead, they are typically displayed with the lowest numbers having the darkest shades of gray and the highest numbers having the lightest shades. The numbers stored in the can actually have any range, but they are usually scaled to the *full dynamic range* of 0–255 before the image is displayed, which we call the screen value. A screen value of 0 displays the color black, and a screen value of 255 is white.

As mentioned before, images are grids of numbers, and, as with any set of numbers, statistics can be used to describe them. We can plot a histogram of the brightness values, as well as calculate an average and a standard deviation. We can also consider the image to be a matrix and perform matrix algebra on the

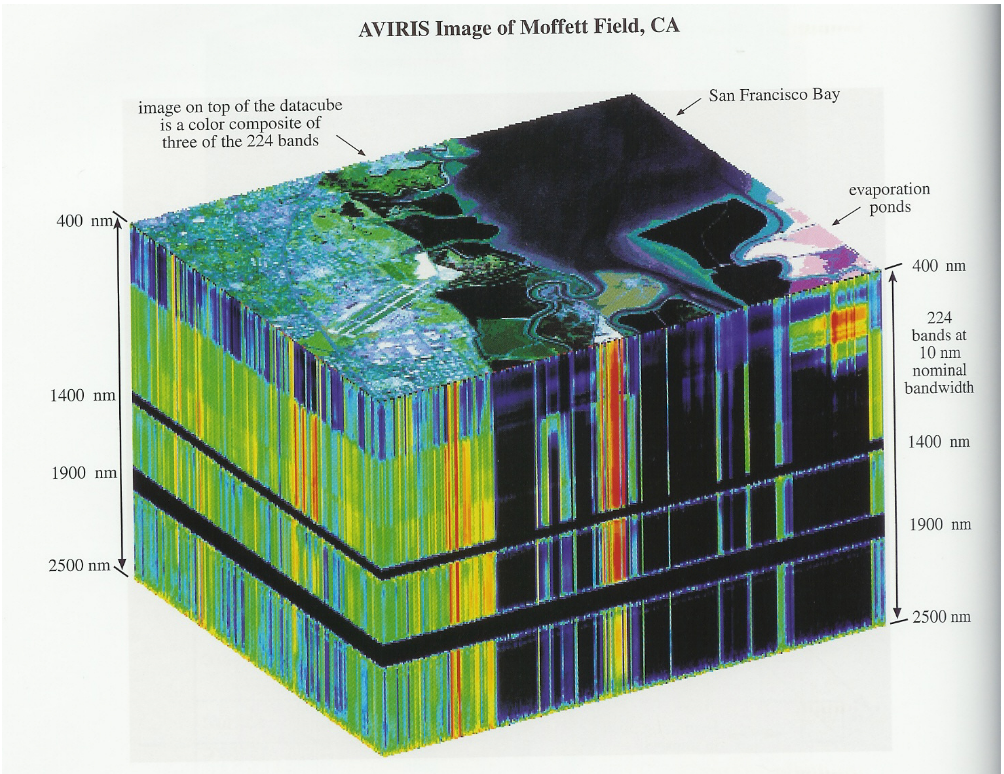

Figure 1: A full AVIRIS image cube. The lateral x- and y-axes correspond to geographic dimensions while the vertical z-axis corresponds to that 2D geographic image at different bands or wavelengths.

image. In fact, many of the higher-order processing techniques we'll learn about later in the semester use matrix algebra. The matrix representing an image will have as many columns and rows as there are pixels in the x- and y-dimensions of the image. The values in each cell of the matrix are equivalent to the brightness values of the corresponding pixel.

To illustrate, let's consider two basic properties of images, *brightness* and *contrast*. Brightness, as its name implies, refers to how bright or dark the overall image is. Contrast refers to how much definition the image has—in images with low contrast, the colors all look the same and it becomes difficult to make out the objects in the image. Figure 2 illustrates these principles. We can begin to quantify these properties (though this is by no means perfect) by realizing that the higher the average of the brightness values, the higher the overall brightness of the image. Furthermore, the higher the standard deviation of the brightness values, the more contrast there is.

What if we want to adjust the brightness or contrast of an image? This is where image math comes in. Recall that images can be thought of as matrices. So, let any image be represented by the matrix $\mathbf{A}$. We can create a new image $\mathbf{B}$ by setting $\mathbf{B} = \mathbf{A} + 2$, which is equivalent to adding the number 2 to every brightness value in the image (subtraction works the same way). If we do this, which property of the image have we changed? We haven't changed the contrast because adding the same number to every pixel does not spread individual brightness values farther apart—in other words, it does not increase the standard deviation of the population of all the brightness values. What we have done, however, is make every pixel's brightness value higher, so we've increased the brightness of the image. To change the contrast, we would need to multiply every pixel's value by 2 (or some other value), since that would cause the numbers to be spread apart (increase the standard deviation).

These basic techniques are not the only way to adjust the contrast of an image, nor are they necessarily the best way. This is simply to introduce you to the basics of image math. In practice, adjusting the contrast of an image is usually done using computer routines, and the technique is usually referred to as *contrast stretching*. ENVI, the software program you will use, has powerful contrast stretching abilities. Note

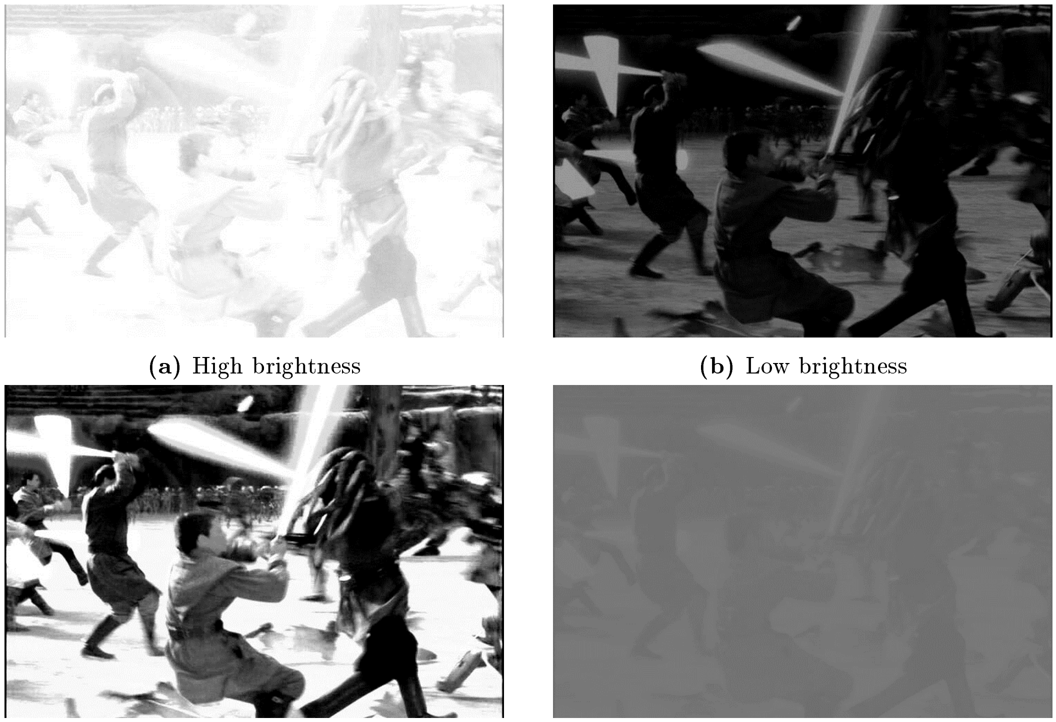

(a) High brightness

(b) Low brightness

(c) High contrast

(d) Low contrast

Figure 2: A grayscale scene from the First Battle of Geonosis from *Star Wars: Episode II - Attack of the Clones* with varying brightness and contrast.

that contrast stretching is, in essence, a manipulation of the image histogram. One very common contrast stretching technique is called *histogram equalization.* This method is good for images with DNs that fall into a narrow range (i.e., have narrow, tall histograms, giving low image contrast). It manipulates the histogram such that it tries to give each pixel value (each bin in the histogram) an equal number of pixels with that value. Figure 3 illustrates this. Figure 3a is a CAT scan of part of a human body with its histogram to the left. Notice the large spikes in the histogram and that the range in pixel values is about 50–150, about half of the total available range of brightness values (i.e., 0–255, the full dynamic range of DN). In Figure 3b, the CAT scan image has been histogram-equalized. The range of the histogram is now 0–255, and each pixel value occurs more frequently. The histogram equalization has brought out more detail in the CAT scan.

What we have discussed so far is manipulation of an image by adding/subtracting or multiplying/dividing an image by a single number. It is also possible to add two images together (although this is complicated if they don't have the same pixel dimensions). In fact, any of the basic mathematical operations can be performed between two images. For example, one image can be multiplied by another, in which case each pixel's value in one image would be multiplied by the corresponding pixel's value in the other image (note that this is <u>not</u> a matrix algebra operation). Why would we do this? As an example, there are special cases of single band images called binary images, in which pixels can have only one of two values—0 or 1. Binary images are often used as image masks. Multiplying some base image by a mask has the effect of keeping some pixels in the base image at their original value (because they are being multiplied by 1), while all others are discarded (because they are being multiplied by 0).

3 Three-Band (Color) Images

We've talked about grayscale images so far, but what about color images? Intuitively, we realize that for an image to be displayed in color, color information must be stored in the image file. How is this accomplished?

From the point of view of our human perception, all colors are mixtures of three primary colors—red,

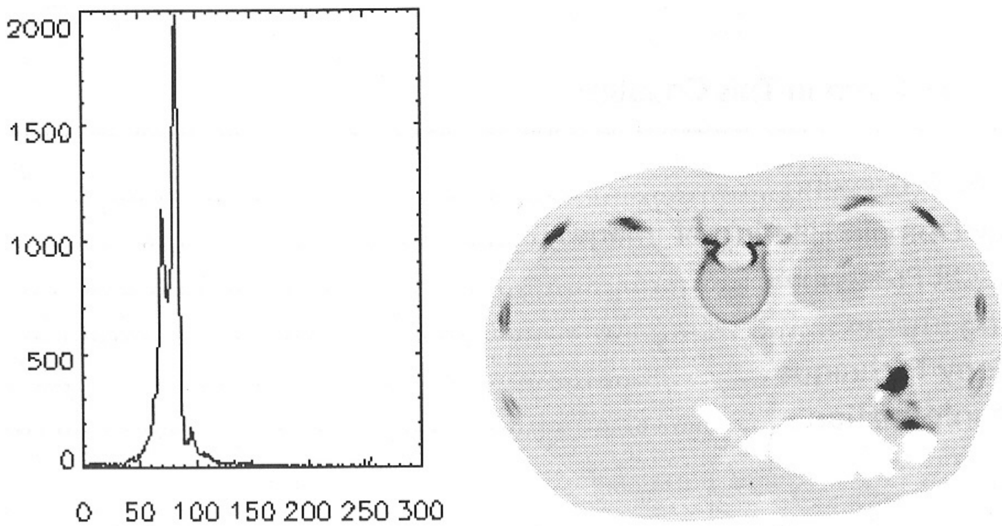

(a) No stretch.

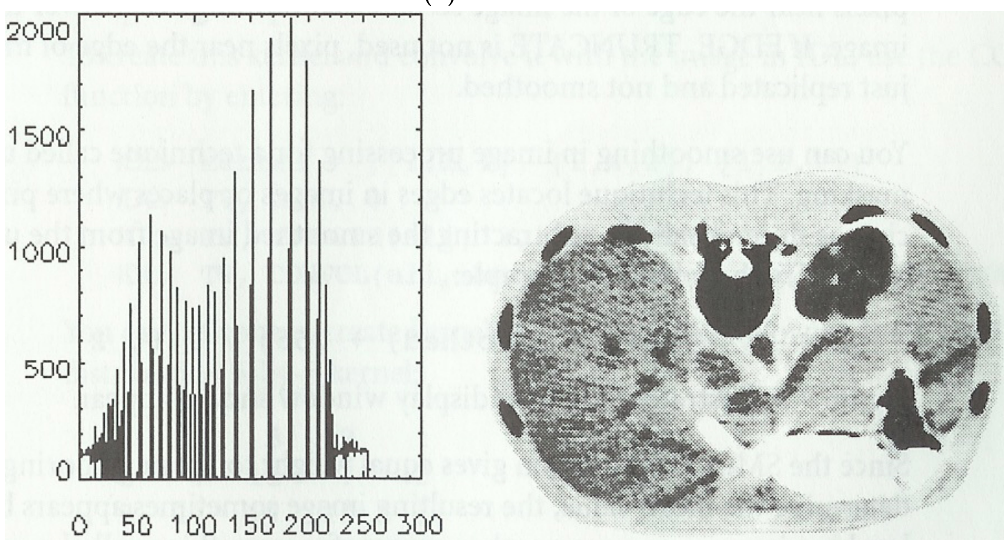

(b) Histogram equalization stretch.

Figure 3: A non-stretched and stretched CAT scan with DN histogram distributions.

green, and blue. (Note: the color mixing we are talking about here is not the same way colors mix when, say, you mix paints. The results are very different.) These 3 primary colors are the basis for color images, since that is the minimum information required to accurately reproduce color in an image. In digital color images, a separate band is created for each of the three primary colors, thus creating a three-band image. These images are often called *RGB images* because of the three colors being represented. The program that displays the image looks at each band's value for each pixel, and, based on their relative proportions, displays the appropriate color on the screen.

It is important to develop an intuitive feel for what proportions of red, green, and blue will produce what color in an image. Recall that when displayed, whatever range the image data falls in is usually scaled to fit the range 0–255. This is true for each band in an RGB image as well, meaning each pixel of an RGB image has a set of three values in this range that determines its overall colors. This set of values is often called a *color triplet* and is always ordered R-G-B. The color black is represented by the triplet [0 0 0] while white is represented by the triplet [255 255 255]. Pure red is [255 0 0], pure green is [0 255 0], and pure blue is [0 0 255]. What happens with equal proportions of red and green but no blue? We get the color yellow! Figure 4 illustrates this *additive color theory* (as opposed to subtractive color theory used with paints and inks).

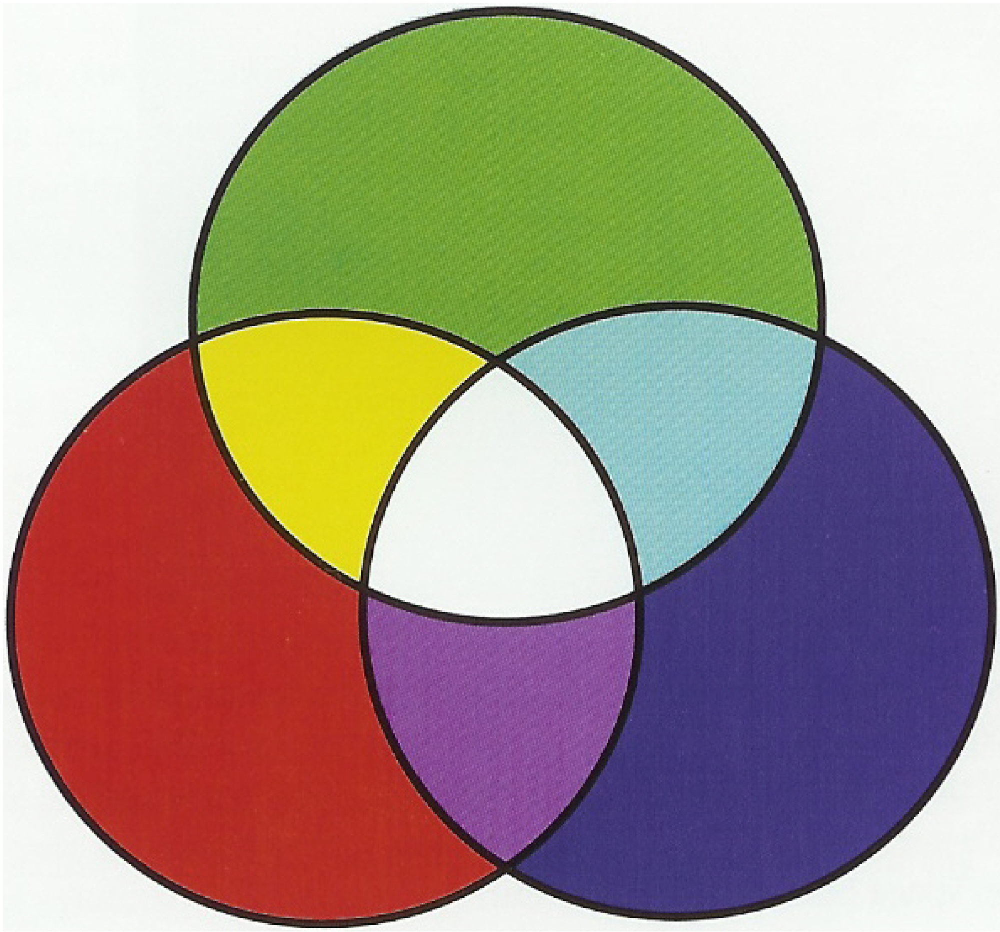

Figure 4: Diagram illustrating additive color mixing.

The contrast stretching that can be done on a single-band image can also be done for each of the individual bands that make up the RGB image (as well as for more complex images cubes with any number of bands). This gives the image processor the power to exert a lot of control over the information the image displays. As you'll learn, frequently noise or some object in the image (like a cloud or a shadow) will complicate your efforts. Contrast stretching is one tool that can be used to overcome the limitations of some of these impairments. Shadowed areas can be stretched, for example, to bring out detail otherwise masked by the shadow. Stretching has the potential to show details that are important in an image while hiding those that are not. Therefore, you the student should learn to stretch the contrast in images in ways that are important. For example, if your image contains very dark portions, a very light portion, and some parts somewhere in between, it can easily be the case that the lightest and darkest parts of the image have no detail because most of the contrast is in the areas of the image with brightness values that are not extreme. Water areas like streams can often show up very dark in images where land is also shown, and in that case it would be very difficult to pick out detail in the stream. If you want to see sediment flux in the river, for example, you should do your stretch so that the pixel values in the darker areas are spread out. In a color image, you have even greater control. If you are working on an image of a red car, the blue and green channels will likely not have a lot of detail in them. You might be able to stretch the contrast in these channels and bring out more detail in the image.

Land Remote Sensing
Lab 2: Basics of Image Processing

Assigned: September 8, 2017
Due: September 22, 2017

If you need more space to explain your answers, then please write on the back of the page or attach additional pages. Within this lab, you will be asked to generate several image products. Please sensibly name these images, collect them into a single .zip file named `xxxxxx_lab02.zip` where xxxxxx is your last name, and e-mail them to `gleggers@gatech.edu` with the subject line of "`[EAS 8803] Lab 02 Products`."

Any files you need are at: http://wray.eas.gatech.edu/remotesensing2017/RS_lab02_files.zip.

1 ENVI Introduction

ENVI (ENvironment for Visualizing Images) is an IDL-based software application by Harris Geospatial used to process and analyze geospatial and spectral imagery. It is the principal tool we will use in this and future labs. ENVI includes many tutorials to help learn its various processes, and you will be asked to complete two introductory ones to familiarize yourself with the software's basic functionality.

Launch the ENVI Classic software. From the *Help* drop-down menu, select *Start ENVI Classic Help*. In the new system page that opens, select *Getting Started* → *Tutorials* → *ENVI Classic*. From here, please complete the *A Quick Start to ENVI Classic* and *Introduction to ENVI Classic* tutorials. The files needed for the tutorials are included in the `tutorial` subdirectory of the provided .zip file.

2 Image Math

1. Use the following "images" the answer the questions below:

A

1	4	2
2	4	1
3	5	1

B

7	5	6
1	2	1
8	6	9

(a) **A** + 2:

(b) **A** + **B**:

(c) Average of **A**:

(d) Standard deviation of **B**

(e) $\mathbf{B} \times 3$

(f) Standard deviation of ($\mathbf{B} \times 3$)

(g) Histogram of $\mathbf{A}$

3 Stretching, Brightness, and Contrast

2. Draw hypothetical but valid histograms for images with the following characteristics:

(a) high contrast, low brightness

(b) moderate brightness, low contrast

(c) moderate brightness, high contrast

(d) low brightness, moderate contrast

(e) moderate brightness, zero contrast

3. In ENVI, open the Landsat single-band image file `LT50130322002275LGS01_B7.tif`, and display it in a new window. View a histogram of the data numbers (DNs), e.g. by clicking *Enhance* → *Interactive Stretching...* in ENVI Classic version.

   (a) What are the minimum and maximum DNs in this image?

   (b) Describe the shape of the histogram. How many peaks does it have, and what are their causes? What is the approximate DN range covered by each peak?

   Without changing ENVI's default contrast/brightness stretch, locate Manhattan Island (centered at roughly 40°46' N, 73°58' W).

   (c) Is the default stretch the most useful way to view land surfaces such as this? Why or why not?

   (d) Create your own preferred stretch, and save it as an image file. Briefly explain how it differs from the ENVI default and why you prefer it. E-mail your image.

4 RGB Triplets and Multispectral Imagery

4. Give reasonable colors for these RGB triplet values:

   (a) [51 255 255] ______________________

   (b) [102 51 0] ______________________

   (c) [255 0 153] ______________________

   (d) [102 102 102] ______________________

5. Give reasonable RGB triplet values for these colors:

   (a) orange ______________________

   (b) black ______________________

   (c) tan ______________________

   (d) cyan ______________________

6. Open the image file `qb_boulder_msi` in ENVI. This is a multispectral image of Boulder, CO from the QuickBird2 satellite owned and operated by DigitalGlobe. Using three of the available bands, create a "true-color" RGB image.

   (a) Which bands did you assign to each of R, G, and B and why? E-mail your image.

   Now locate this area in Google Earth.

   (b) What is the name of the irregular dark area in the right half of the image?

In a new display window, set the 830 nm, 660 nm, and 560 nm data as R, G, and B, respectively. Link this window to your true-color display. Some areas in the new display window have a bright red color.

(c) What are these bright red areas, and why do they have this spectral character?

5 Examples of Image Manipulation

A common use of remote sensing is *change detection*. In this application, two (or more) images of the same area are taken at different times and compared. Can you figure out a simple way to do this using band math techniques you have learned so far? You are provided two images, `cover2002.tif` and `cover2003.tif`, that show a simple hypothetical coastline in 2002 and 2003, respectively.

7. Produce an image that shows the amount of coast line that eroded between 2002 and 2003. Describe your process. If these were real images, what would be some problems that could occur (or problems of which you should be cognizant) if you used your devised technique. How would they make your technique less effective? E-mail your image.

Digital elevation models (DEMs) are images where pixel values represent topography, so bright colors are high elevation and dark colors are lower elevation. The file `dem.img` represents a DEM with elevation in feet.

8. Suppose you want to conduct a study only in areas of certain elevations.

   (a) Make an image in which all elevations above 1300 ft are preserved, and all values below that are set to 0. Describe your process, and e-mail your image.

(b) Also, make an image that only shows elevations between 1250 and 1275 ft. Describe your process, and e-mail your image.

Computer graphics (CG)-generated special effects are a major modern application of image processing, particularly in blockbuster Hollywood movies.

9. Using image math and the file `mask.img`, insert the figure of Mace Windu in `mace.jpg` into `background.jpg`. Describe your process, and e-mail your image.